

# Born to Learn: the Inspiration, Progress, and Future of Evolved Plastic Artificial Neural Networks

Andrea Soltoggio, Kenneth O. Stanley, Sebastian Risi

**Abstract**—Biological neural networks are systems of extraordinary computational capabilities shaped by evolution, development, and lifetime learning. The interplay of these elements leads to the emergence of adaptive behavior and intelligence, but the complexity of the whole system of interactions is an obstacle to the understanding of the key factors at play. Inspired by such intricate natural phenomena, Evolved Plastic Artificial Neural Networks (EPANNs) use simulated evolution in-silico to breed plastic neural networks, artificial systems composed of sensors, outputs, and plastic components that change in response to sensory-output experiences in an environment. These systems may reveal key algorithmic ingredients of adaptation, autonomously discover novel adaptive algorithms, and lead to hypotheses on the emergence of biological adaptation. EPANNs have seen considerable progress over the last two decades. Current scientific and technological advances in artificial neural networks are now setting the conditions for radically new approaches and results. In particular, the limitations of hand-designed structures and algorithms currently used in most deep neural networks could be overcome by more flexible and innovative solutions. This paper brings together a variety of inspiring ideas that define the field of EPANNs. The main computational methods and results are reviewed. Finally, new opportunities and developments are presented.

**Index Terms**—Artificial Neural Networks, Learning, Plasticity, Evolution.

## I. INTRODUCTION

Over the course of millions of years, evolution has led to the emergence of innumerable biological systems, and intelligence itself, crowned by the discovery of the human brain. Evolution, development, and learning are the fundamental processes that underpin biological intelligence. Thus, it is no surprise that scientists have modeled artificial systems to reproduce such phenomena (Sanchez et al., 1996; Sipper et al., 1997; Dawkins, 2003). However, our current knowledge of evolution, biology, and neuroscience remains insufficient to provide clear guidance on the essential mechanisms that are key to the emergence of such complex systems.

Fields such as artificial intelligence and artificial life (Langton, 1997) address questions about the emergence of life and intelligence from a computational perspective, thus abstracting the principles from the medium, i.e. biological or in-silico, and attempting to find algorithms that reproduce properties of their biological counterparts. The simulation of natural evolution in-silico is a tool that, besides being an optimization

algorithm (Bäck and Schwefel, 1993), attempts to discover pivotal elements to artificially evolve life and intelligence. Thus, the field of evolutionary computation (Holland, 1975; Eiben and Smith, 2015) includes areas such as evolutionary robotics (Harvey et al., 1997; Nolfi and Floreano, 2000) and neuroevolution (Yao, 1999) to investigate the autonomous design of complex, bio-inspired adaptive physical and neural systems.

This paper frames the field that attempts to *evolve plastic artificial neural networks*, and introduces the acronym EPANN. EPANNs are *evolved* because parts of their design are determined by an evolutionary algorithm; they are *plastic* because they undergo various time-scale changes, beyond neural activity, while experiencing sensory-motor information streams during a lifetime simulation. The final capabilities of such networks are a result of genetic instructions, determined by evolution, that enable learning once the network is placed in an environment. Static ANNs with evolved connection weights (Yao, 1999) are not considered EPANN.

The field studying EPANNs is motivated by the long-term ambitious promise of discovering principles of neural adaptation, learning, and memory, an endeavor that also entails a number of challenges. A primary challenge is that the interactions of dynamics across the evolutionary and learning dimensions make the whole system difficult to design and analyze. Moreover, it is not clear what key aspects are necessary and what abstractions can be made to achieve the objective. A second problem is the high computational cost required to simulate even simple models of lifetime learning and evolution. Thirdly, open-ended evolutionary experiments are difficult to benchmark because they are based on loosely defined objectives such as increasing behavioral complexity, intelligence, adaptability, and evolvability (Miconi, 2008). For this reason, EPANNs rarely compete with ad-hoc machine learning algorithms specifically designed to improve performance on well-defined and narrow problems. The lack of benchmarks and commercially viable applications make it difficult to assess the scientific value or the pathways to impact for many open-ended evolutionary systems. However, in this paper we argue that the above challenges can reveal untapped directions of scientific investigation with a high potential for discovery. Advances in this field have improved our understanding of artificial and biological intelligence and will continue to do so.

EPANNs rely on two autonomous search processes: evolution and learning, which arguably place them among the most advanced AI and machine learning systems in terms of open-endedness and human-free design. Those systems rely

Department of Computer Science, Loughborough University, LE11 3TU, Loughborough, UK, a.soltoggio@lboro.ac.uk

Department of Computer Science, University of Central Florida, Orlando, FL, USA, kstanley@cs.ucf.edu

IT University of Copenhagen, Copenhagen, Denmark, sebr@itu.dk

the least on pre-programmed instructions, and the most on the autonomous discovery of intelligence and adaptation from experience and from the interaction with a real or simulated world.

In recent years, progress in a number of relevant areas have set the stage for renewed advancement of EPANNs: more focus has been devoted to learning networks and deep neural networks; there has been a remarkable increase in computational power by means of parallel GPU computing; theories on search, complexity, and genotype-phenotype mapping now allow for less naive evolutionary approaches; high-bandwidth sensors, actuators, dedicated hardware, and big data are becoming increasingly more available; new hardware paradigms such as neuromorphic chips are being developed; and finally, neuroscience and genetics provide us with an increasingly large set of inspirational principles. This progress has changed the theoretical and technological landscape in which EPANNs first emerged, providing greater research opportunities than in the past.

This paper brings together the principles that inspire and justify research in EPANNs to delineate a unique research field (Section II), it reviews the progress so far (Section III), and it outlines current and future challenges and opportunities (Section IV).

## II. INSPIRATION

EPANNs are strongly inspired by the biological evidence of natural computation and intelligence (Floreano and Mattiussi, 2008; Downing, 2015). However, the notion of “nature-inspired” extends to a wide breadth of knowledge, is often loose, and is subject to different interpretations. This section aims to pinpoint the main ideas that motivate research specifically in EPANNs.

Biological neural systems, the primary inspiration for EPANNs, are complex computational structures responsible for a variety of functions from perception, to motor skills, to behavior. In neuroscience, an important question is, *what makes a brain do what it does?* In the computational field that includes EPANNs, more relevant is to understand *what are the fundamental ingredients and design principles discovered by nature that resulted in high levels of adaptation and intelligence.* Those principles might then be abstracted and implemented in original ways in-silico to achieve similar results to those observed in nature. Crucially, artificial systems do not need to implement the same constraints and limitation as biological systems (Bullinaria, 2003). Thus, inspiration is not simple imitation.

We now know that genetic instructions determine the ultimate capabilities of a brain (Deary et al., 2009; Hopkins et al., 2014). Different animal species manifest different levels of skills and intelligence because of their genetic blueprint, given by the phylogenetic path of a particular species (Schreiweis et al., 2014). We also know that the vastly intricate structure of the brain emerges from one single zygote cell through a developmental process (Kolb and Gibb, 2011). Finally, the developmental process is strongly shaped by input-output learning experiences throughout early life (Hensch et al., 1998;

Kolb and Gibb, 2011), and yet high levels of plasticity are maintained throughout the entire lifespan (Merzenich et al., 1984; Kiyota, 2017).

These dimensions, *evolution, development and learning*, also known as the phylogenetic (evolution), ontogenetic (development) and epigenetic (learning) (POE) dimensions (Sipper et al., 1997) are the pillars of biological plastic brains. This review focuses on evolution and learning, but less on artificial developmental systems (Stanley and Miikkulainen, 2003b) because development may occur independently of, or dependent upon, external factors. In models where it occurs independently of external factors, development is entirely driven by genetic instructions, and does not include learning; in the second case, when development is affected by external factors, part of the process can be interpreted as learning. Thus, learning in EPANNs may occur by means of a variety of plasticity mechanisms, some of which could be seen as developmental processes.

In the design of bio-inspired artificial intelligence, an interesting issue is whether all three POE dimensions are equally important, or even necessary at all. Opinions are diverse, in particular in relation to *what key aspects produce adaptation and intelligence.* Recent advances in neuroscience suggest that plasticity in the brain is far more responsible for the manifestation of behaviors and skills than previously thought (LeDoux, 2003; Doidge, 2007; Grossberg, 2012), meaning that behavior is less determined by genetics and more by stimuli-driven plasticity. A variety of social behaviors are learned with experience in different circumstances. Skills such as reading, playing a musical instrument, or driving a car, are mastered even if none of those behaviors existed during evolutionary time. Thus, genetic instructions are responsible mainly for setting recipes to learn.

Recent developments in a different field, that of machine learning (Michalski et al., 2013; Alpaydin, 2014) and neural learning (LeCun et al., 2015), also suggest the importance of learning from large input-output data. Deep reinforcement learning has demonstrated a remarkable capability to learn complex input-output relations using extensive training (Mnih et al., 2015; Silver et al., 2016). Deep learning has pioneered a remarkable progress in image classification and speech recognition (Schmidhuber, 2015; Deng et al., 2013). Other areas of cognition such as the capabilities to make predictions (Hawkins and Blakeslee, 2007), establish associations (Rescorla, 2014) and regulate behaviors (Carver and Scheier, 2012) are also based on learning from experience but have seen a slower progress in computational methods.

It is commonly agreed that EPANN’ final objective is not to perform well on one specific task, but rather to perform well on adaptation processes such as learning the structure of a particular world, its associations, patterns, and possible behaviors. In this endeavor, the evolutionary computation community, and particularly the field that studies neuroevolution, is extending research efforts to include open-ended evolutionary environments in which many tasks and behaviors might be required (Miconi, 2008). In this way, both biological observations and computational tools are yielding inspirational principles for EPANNs.

Whilst the range of inspiring ideas is large and heterogeneous, the analysis in this review proposes that EPANNs build upon five main concepts:

- *evolutionary processes*,
- *inspiration from biological neural networks*,
- *abstractions in brain simulations*,
- *artificial plastic neural networks*, and
- *intelligence-testing environments*.

These broad categories encompass a number of research areas that range from large fields such as brain physiology to specific themes such as artificial genetic encoding. Figure 1 illustrates the five main concepts and the relevant research areas to give a graphical representation of the variety and overlap of ideas that contribute to EPANNs. The sections that follow describe each of these main concepts.

#### A. Evolutionary processes

A central idea in evolutionary computation (Goldberg and Holland, 1988; Back et al., 1997) is that evolutionary processes, i.e. those that occurred in nature during the course of billions of years, can also be simulated with computer software. This idea led to the belief that intelligent computer programs could emerge with little human intervention by means of evolution in-silico (Holland and Reitman, 1977; Koza, 1992; Fogel, 2006).

The emergence of intelligent evolved software, however, did not occur as easily as initially hoped. The reasons for the slow progress are not completely understood, but a number of problems have been identified in relation to levels of abstraction, diversity, selection criteria, fitness and simulation environments, genotype-phenotype mapping, and available computation power. Moreover, in the biological evolution of intelligent and complex forms of life, it is not clear yet which stepping stones were more challenging for natural evolution to overcome (Stanley and Lehman, 2015). Similarly, current evolutionary algorithms might not yet focus on the pivotal search aspects that are necessary to evolve intelligence.

To address these questions, research topics in evolutionary computation, particularly those that contribute to EPANNs, investigate the concepts of evolvability and scalability (Wagner and Altenberg, 1996; Pigliucci, 2008; Lehman and Stanley, 2013), the encoding of genetic information through the genotype-phenotype mapping processes (Hornby et al., 2002), and novelty search and the deception of fitness objectives (Lehman and Stanley, 2008; Stanley and Lehman, 2015). Popular algorithms applied to EPANNs are reviewed later in Section III-A.

#### B. Inspiration from biological neural networks

Biological networks are a source of inspiration for their complexity and the variety of phenomena they produce. Lifetime learning includes astonishing skills such as social behavior, playing an instrument, and speaking one or more languages. Many of the skills we possess are acquired by experiencing stimuli and actions.

The fact that experiences guide life-long learning is an old and intuitive notion. Scientific evidence of animal learning

from experience was documented initially in the works of *behaviorism* with scientists such as Thorndike (1911), Pavlov (1927), Skinner (1938, 1953), and Hull (1943). The main idea was to experiment and verify scientifically how experiences cause a change in behavior, in particular as a result of learning associations and observable behavioral patterns (Staddon, 1983). This approach means linking behavior to brain mechanisms and dynamics, an idea initially entertained by Freud (K $\phi$ PPE, 1983) and later by other illustrious scientists trying to bring physiology and psychology together (Hebb, 1949; Kandel, 2007). A seminal contribution to link physiology to psychology comes from Hebb (1949), whose ubiquitous principle that neurons that *fire together, wire together* is relevant to both low level neural wiring and high level psychological theories (Doidge, 2007).

The seminal work of Kandel and Tauc (1965), and following studies (Clark and Kandel, 1984), were the first to demonstrate that changes in the strength of connectivity among neurons related to behavior learning. Walters and Byrne (1983) showed that a single neuron can perform associative learning such as classical conditioning, a class of learning that is observed in simple neural systems such as by the Aplysia (Carew et al., 1981). Plasticity driven by local neural stimuli, i.e. compatible with the Hebb synapse (Hebb, 1949), is responsible not only for fine tuning, but also for building a working visual system in the cat's visual cortex (Rauschecker and Singer, 1981). Those advances in the 1970s and 1980s lead to the mathematical formulation of synaptic learning dynamics (Lisman, 1989; Brown et al., 1990), e.g. the BCM rule, inspired by the cat's visual cortex (Bienenstock et al., 1982), and the Oja's rule (Oja, 1982). A Hebbian-compatible rule that regulates synaptic changes according to the firing times of the presynaptic and postsynaptic neurons was observed by Markram et al. (1997) and named Spike-Timing-Dependent Plasticity (STDP).

Besides fine tuning of synaptic connections, a plastic brain is capable of establishing new pathways and connections by means of what is referred to as structural plasticity (Lamprecht and LeDoux, 2004; Pascual-Leone et al., 2005; Russo et al., 2010). Structural plasticity occurs primarily by means of processes taking place during development. Axon growth is known to be regulated by neurotrophic nerve growth factors that guide axons to extend in particular directions (Tessier-Lavigne and Goodman, 1996). Developmental processes and neural plasticity are often indistinguishable (Kolb, 1989; Pascual-Leone et al., 2005) because the brain is highly plastic during development. There is evidence that structural plasticity continues well into adulthood (Pascual-Leone et al., 2005). Our understanding of structural plasticity in shaping both thoughts and behavior has grown significantly. Neuroscientific advances reviewed in Damasio (1999); LeDoux (2003); Pascual-Leone et al. (2005); Doidge (2007); Draganski and May (2008) reveal that fundamental sensations, motor patterns, associations, and ways of thinking are less instinctive and hardwired than previously thought. Both structural and functional plasticity in biology is essential to acquiring long-lasting new skills, and for this reason appears to be an important point of inspiration for EPANN.

One further plasticity mechanism that has attracted attention

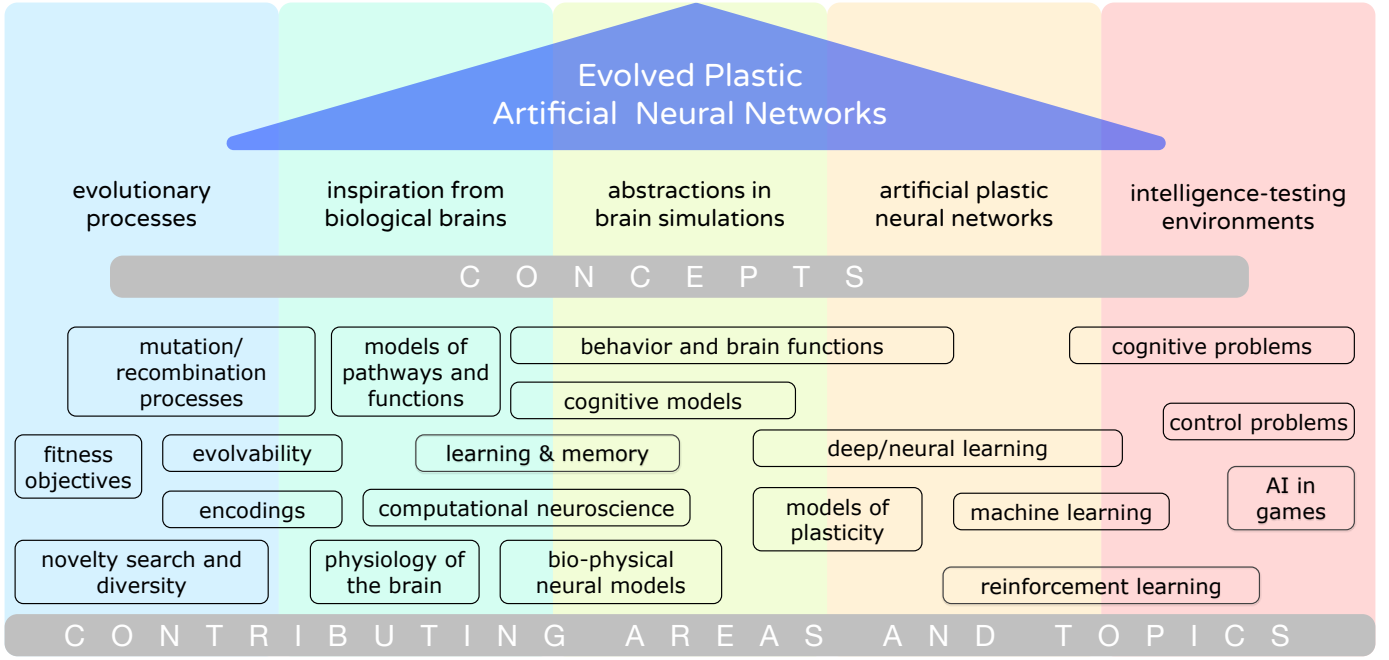


Figure 1: The five main concepts that contribute to EPANNs draw from diverse research areas and topics.

is *neuromodulation*, or *heterosynaptic plasticity* (Marder and Thirumalai, 2002; Gu, 2002; Bailey et al., 2000). Neuromodulation appears to mediate plasticity by increasing or decreasing its effect according to a modulatory signal. A number of modulatory chemicals such as acetylcholine (ACh), norepinephrine (NE), serotonin (5-HT) and dopamine (DA) appear to regulate a large variety of neural functions, from arousal and behavior (Harris-Warrick and Marder, 1991; Hasselmo and Schnell, 1994; Marder, 1996; Katz, 1995; Katz and Frost, 1996), to pattern generation (Katz et al., 1994), to memory consolidation (Kupfermann, 1987; Hasselmo, 1995; Marder, 1996; Hasselmo, 1999). Learning by reward in monkeys was linked to dopaminergic activity during the 1990s with studies by Schultz et al. (1993, 1997); Schultz (1998). For these reasons, neuromodulation is considered an essential element in cognitive and behavioral processes, and has been the topic of a considerable amount of work in EPANN (Section III-E).

In short, plasticity in the brain appears driven by an extraordinarily rich set of dynamics and chemicals. From these biological observations, computational questions relevant to EPANN include: (1) How does the brain structure grow – driven both by genetic instructions and neural activity – to acquire functions and behavior? (2) What are the key plasticity mechanisms from biology that can be applied to EPANN? (3) Are memories, skills, and behaviors stored in plastic synaptic connections, in patterns of activities, or a combination of both? Whilst neuroscience continues to provide inspiration and insight into plasticity in biological brains, EPANNs serve the complementary objective of implementing, seeking, and verifying designs of bio-inspired methods for adaptation, learning, and intelligent behavior.

### C. Abstractions in brain simulations

If an accurate brain simulator existed, the model would manifest most, if not all, the behavioral and cognitive properties of a biological brain. Thus, the attempt to implement intelligence by simulating a complex biological network of neurons appears justified. However, despite considerable progress, current neural models are far from exhibiting sophisticated cognitive abilities. It is thus a legitimate question to ask what is missing from current models of the brain to match animal and human intelligence.

Neural models range from small directed graphs in which nodes are simple functions of the input, to complex and sometimes large spiking neural models and bio-physical models (Markram, 2006, 2012; Eliasmith et al., 2012). The variety and range of neural models derive from their diverse aims, but also from the current uncertainty on the level of abstraction required to reproduce those behaviors typical of human and animal intelligence.

Whilst modeling neural cells and their connectivity is one approach to mimicking brain function, a second approach is based on information processing and statistics. Computational neuroscience seeks to make sense of the computation in the brain from a theoretical view point, thus understanding neural dynamics, processes and computation from a mathematical and statistical perspective. Computational neuroscience is concerned primarily with biological brains, so adhering to biological properties and measurements is of particular importance. Although brain models and computational neuroscience do not seek to create EPANNs, they can provide tools to design such systems.

### D. Artificial plastic neural networks

Artificial neural networks are plastic whenever a learning algorithm is applied to modify their properties such as the

topology or the connectivity. Because networks perform correctly only with specific weight configurations, learning rules are an integral part of most neural networks; see Widrow and Lehr (1990) for an early review. One exception, not discussed in this paper, is when evolution is used to search static connection weights of a network, i.e. as a surrogate of learning rules (Yao, 1999).

EPANNs are networks that, as opposed to hand-designed ANNs, employ evolution as a meta-learning process in which the learning mechanisms are evolved. Testing occurs with behavioral simulations that often apply online learning, i.e. the learning is active at each time step. Established learning rules are often used as *ingredients* or inspiration in EPANNs, which then search for good parameter configurations, efficient combinations of rules and network topologies, new functions representing novel learning rules, etc. EPANN experiments are suited to include the largest possible variety of rules because of (1) the variety of possible tasks in a simulated behavioral experiment and (2) the flexibility of evolution to combine rules with no assumptions about their dynamics. The review that follows gives a snapshot of the extent and scope of various learning algorithms for ANN that can effectively be building blocks of EPANN.

In supervised learning, *backpropagation* is the most popular learning rule used to train both shallow and deep networks (Rumelhart et al., 1988; LeCun et al., 2015), typically for classification in feed-forward networks. Unsupervised learning is implemented in neural networks with self-Organising Maps (SOM) (Kohonen, 1982, 1990), auto-encoders (Bourlard and Kamp, 1988), restricted Boltzmann machines (RBM) (Hinton and Salakhutdinov, 2006), and Hebbian plasticity (Hebb, 1949; Gerstner and Kistler, 2002a; Cooper, 2005). Hebbian rules, in particular, given their biological plausibility and unsupervised flavor of learning, are a particularly important inspirational principle. Variations have been proposed to include, e.g. terms to achieve stability (Oja, 1982; Bienenstock et al., 1982) and various constraints (Miller and Mackay, 1994), or more advanced update dynamics such as dual weights for fast and slow decay (Levy and Bairaktaris, 1995; Soltoggio, 2015). Those rules depend often on the precise topology of the network and are hand-designed, although cost-minimization Hebbian rules have been proposed recently (Pehlevan et al., 2015; Bahroun and Soltoggio, 2017).

Often, behavior learning requires a measure of the quality of such behaviors, implemented as positive or negative values that represent the consequences of stimulus-action sequences. This type of learning is also referred to as trial and error, or operant reward learning in animal learning, and was formalized mathematically as Reinforcement Learning (RL) (Sutton and Barto, 1998). RL in neural networks (Pennartz, 1997) makes use of robotic or agent-environment problems with a reward signal that affects plasticity in the network (Lin, 1993; Niv et al., 2002; Soltoggio and Stanley, 2012; Auerbach et al., 2014; Pugh et al., 2014).

Neuromodulated plasticity (Fellous and Linster, 1998) is an important mechanism in EPANNs (reviewed later in Section III-E). Modulation of signals and gated learning (Abbott, 1990) is inspired by the idea that behavior learning and reg-

ulation of behavior require neuromodulatory activity (Baxter et al., 1999; Suri et al., 2001; Birmingham, 2001; Alexander and Sporns, 2002; Doya, 2002; Fujii et al., 2002; Suri, 2002; Ziemke and Thieme, 2002; Sporns and Alexander, 2003).

Plastic neural models are used also to demonstrate how behavior can emerge from a particular circuitry modeled after biological brains. Computational models of, e.g., the basal ganglia and modulatory systems may propose plasticity mechanisms and aim to demonstrate the computational relations among various nuclei, pathways, and learning processes (Krichmar, 2008; Vitay and Hamker, 2010; Schroll and Hamker, 2015).

Finally, plasticity rules for spiking neural networks (Maass and Bishop, 2001) aim to demonstrate unique learning mechanisms that emerge from spiking dynamics (Izhikevich, 2007), as well as model biological synaptic plasticity (Gerstner and Kistler, 2002b).

Plasticity in neural networks, when continuously active, was also observed to cause catastrophic forgetting (Robins, 1995). If learning occurs continuously, i.e. online, new information or skills have the potential to overwrite previously acquired information or skills (Abraham and Robins, 2005). Metaplasticity and consolidation mechanisms were suggested to address this problem (Finnie and Nader, 2012; Soltoggio, 2015).

In conclusion, plasticity is central to most approaches to ANNs. The variety of algorithms indicate that different approaches were developed to solve different problems. In this context, EPANNs offer a unique tool to test hypotheses on the effectiveness and suitability of different models, or their combination, in a variety of different scenarios.

### E. Intelligence-testing environments

Evolution works on organisms within an environment to exploit its resources, find survival solutions, and adapt to changes (Darwin, 1859). However, what makes an environment conducive to the evolution of learning is not clear. In particular, what are the challenges faced by learning organisms in the natural world, and how can those be abstracted and ported to a virtual, simulated environment?

In the early phases of AI, logic and reasoning were thought to be key to intelligence (Cervier, 1993), so symbolic input-output mappings were employed as tests. Soon it became evident that intelligence is not only symbol manipulation, but resides also in subsymbolic problem solving abilities emerging from the interaction of brain, body and environment (Steels, 1993; Sims, 1994), so more complex simulators of real-life environment were developed. A further step in complexity is represented by high-level planning and strategies required, e.g., when applying AI to games (Allis et al., 1994; Millington and Funge, 2016) or articulated robotic tasks. Planning and decision making with high bandwidth sensory-motor information flow such as those required for humanoid robots or self-driving vehicles are current benchmarks of intelligent systems. Finally, environments in which affective dynamics and feelings play a role are recognized as important for human well being (De Botton, 2016). Detecting human emotions (Lee and Narayanan, 2005) and providing emotional support will

increase in importance as a benchmark for AI. The increase in the variety and complexity of simulated environments offers new opportunities to develop EPANN as further outlined later in Section 4, on open challenges.

### III. PROGRESS ON EVOLVING ARTIFICIAL PLASTIC NEURAL NETWORKS

EPANNs may draw inspiration from one or more theories and techniques overviewed in the previous section. Specific research questions under investigation determine the sets of chosen algorithms. One common focus is using evolution to search for emergent mechanisms of adaptation, intelligence, and complexity. Thus, the objectives may not be only to achieve quality solutions, but primarily to investigate the search path, the evolutionary steps, the features of the final solutions, and other dynamical aspects of the evolutionary process. The hope is that counter-intuitive, distributed, sub-symbolic and minimal algorithms will emerge from a combination of the right building blocks and initial conditions. Within this general idea of searching for emergent mechanisms of learning and adaptation, particular areas of focus, reviewed in the next sections, have contributed to progress with EPANNs.

EPANN setups can be viewed as the interaction of simulated evolution and intelligent testing environments in which the various inspirational principles illustrated in this section are used to enrich the learning dynamics. As simplified graphical representation is given in Fig. 2.

#### A. Evolutionary algorithms for EPANN

When applied to EPANNs, artificial evolution is characterized by systems of growing complexity and changing search spaces (Holland, 1975; Michalewicz, 1994). As such, it differs from parameter optimization (Bäck and Schwefel, 1993) in which search spaces are often static. The growing and changing fitness landscapes affect the choice of the particular evolutionary algorithm to use to evolve plastic networks. Desired features for an evolutionary algorithm to evolve plastic networks are:

- variable genotype length to allow for complexity to grow through evolution;
- indirect genotype to phenotype encodings to enable efficient and evolvable representations;
- mutations and/or recombination of genomes that allow exploration of new skills while maintaining basic functional properties of the solutions;
- expressing regularity, repetition and patterns, which appear to be features in biological structures
- flexibility to operate at different levels of abstraction, e.g. network, neuron, or synaptic level, to allow evolution to extend the search across those levels;
- maintaining diversity in the population to take advantage of parallel search along diverse phylogenetic paths, and to allow increased exploratory search along paths of sub-optimal fitness values;
- devising fitness as an indication of general adaptive or intelligent behavior to provide an evolutionary advantage to solutions expressing those properties.

All points, except the last two, deal with the capability of the algorithm to mimic various biological phenomena within the evolving solution. The last two points deal with the modality of selection, how and which solutions survive.

Many evolutionary algorithms include one or more of those mechanisms (Fogel, 2006). Due to the complexity of neural network design, the field of *neuroevolution* was the first to explore most extensively those mechanisms. Popular algorithms include early work of Angeline et al. (1994) and Yao and Liu (1997) to evolve neural networks, Neuroevolution of Augmenting Topologies (NEAT) (Stanley and Miikkulainen, 2002), Analog Genetic Encoding (AGE) (Mattiussi and Floreano, 2007), and HyperNEAT (Stanley et al., 2009). While these algorithms were not initially devised to evolve plastic networks, they are often adapted to operate with EPANNs.

#### B. Evolution of plasticity rules

Early experiments in EPANN started with the attempt to evolve learning rules to change the connection strength among nodes in a connected graph, and thereby enable online lifetime learning. In contrast to back-propagation of errors (Widrow and Lehr, 1990), considered biologically implausible, the Hebbian principle (Hebb, 1949; Cooper, 2005) is a more popular choice in EPANN because is biologically plausible due its dependency solely on local activity values. In its simplest form, the synaptic strength  $w$  is updated by the product of presynaptic ( $x$ ) and postsynaptic ( $y$ ) activities, and a learning rate  $\eta$ :

$$\Delta w = \eta \cdot x \cdot y \quad . \quad (1)$$

More generally, the synaptic change can be expressed as an arbitrary function of local values:

$$\Delta w = f(x, y, w, \theta) \quad , \quad (2)$$

where  $\theta$  are parameters of the function. The type of function and its parameters can be searched by evolution.

Bengio et al. (1990, 1992) proposed the optimisation of a learning rule via evolutionary search. The idea was to use evolutionary search to optimize parameters according to an error (or cost) function. Those studies are also among the first to include a modulatory term in the learning rules. Chalmers (1990) evolved a global learning rule (i.e. a rule that applies to every connection) and discovered that the evolved rule was similar to the well-known delta rule, or Widrow-Hoff rule (Widrow et al., 1960), used in backpropagation. Fontanari and Meir (1991) build on Chalmers (1990)'s approach to evolve a learning algorithm for single-layer networks with binary weights. Baxter (1992) evolved a local learning rule and the architecture for a network that learned four boolean functions of one variable. The network was called Local Binary Neural Net (LBNNs): each connection had a rule that took local neural values and changed its weight to one of two possible values ( $\pm 1$ ). The ideas of evolving learning was implemented in Nolfi and Parisi (1993) by evolving "auto-teaching" inputs, which could then provide an error signal for the network to adjust weights during lifetime. McQuesten and Miikkulainen (1997) showed that neuroevolution can benefit from parent networks teaching their offspring through backpropagation. In



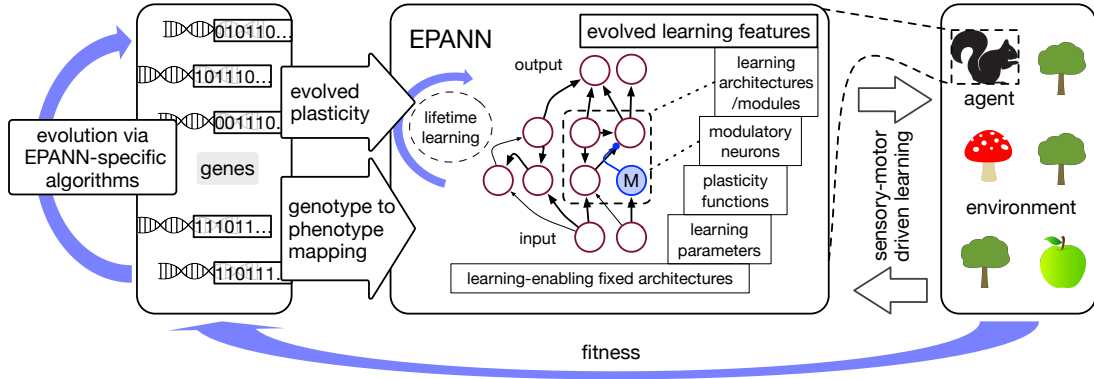


Figure 2: Main elements of an EPANN set-up in which simulated evolution (left box) and an environment (right box) allow for plastic networks to evolve through generations and learn within a lifetime in the environment (central box). The network may include diverse elements, e.g. modules or architectures, modulatory neurons, inputs/outputs, learning parameters, and other structures that enable learning.

Husbands et al. (1998) and Smith et al. (2002), the authors propose a plastic network that adapts its activation functions by simulating the diffusion of a gas. Although the tasks did not involve learning, the network appears to evolve faster than a network evolved with fixed weights.

Such studies that *evolve a network to learn* seek dynamics that can be viewed as a meta-learning process. Abraham (2004) proposed a method called Meta-Learning Evolutionary Artificial Neural Networks (MLEANN) in which evolution searches for initial weights, neural architectures and transfer functions, and finally optimizes existing learning rules. Orchard and Wang (2016) encode the learning rule itself as a network that is then evolved. One general purpose in those studies is to find learning rules, via evolution, that permit a consistent, reliable, and scalable learning in specific problems (Soltoggio, 2008b; Risi and Stanley, 2012).

### C. EPANNs in Evolutionary Robotics

Evolutionary Robotics (ER) brought the idea of autonomously designing robots and controllers by means of evolution (Cliff et al., 1993; Floreano and Mondada, 1994, 1996; Urzelai and Floreano, 2000; Floreano and Nolfi, 2004). One important principle is the idea of *evolving brains with body and environment* (Chiel and Beer, 1997), in which intelligence is embodied because it emerges from the interaction of brain-body-environment (Floreano et al., 2004).

Whilst ER originated with the idea of evolving simulated and real robots with no specific assumptions on neural systems or plasticity (Smith, 2002), experiments suggested that neural control structures evolved with fixed weights perform less well than plastic networks in robotic applications (Nolfi and Parisi, 1996; Floreano and Urzelai, 2001a). In particular, plastic networks were shown to adapt better in the transition from simulation to real robots. This made sense as well as the intuitive concept that evolved plastic neural networks increase the adaptability to changing ER environments (Nolfi and Parisi, 1996). However, the precise nature and the magnitude of those changes that plasticity can cope with is not always easy to quantify in robotic settings. Another study reported

that networks evolved faster, i.e., had better evolvability, when synaptic plasticity and neural architectures are evolved simultaneously (Floreano and Urzelai, 2001b).

The behavioral changes required to switch behaviors in simple ER experiments can also take place with non-plastic recurrent neural networks because inputs can switch the network’s activity to different attractors. These dynamics were observed in particular in recurrent networks in which neurons are modeled as leaky integrators (Beer and Gallagher, 1992; Yamauchi and Beer, 1994), which implement different time constants across neurons. Such networks were called continuous time recurrent neural networks (CTRNN). A study in 2003 observed similar performance in learning when comparing plastic and non-plastic neural networks (Stanley and Miikkulainen, 2003a). That study put forward the idea that the learning of plastic networks as described in Blynell and Floreano (2002, 2003) was at that point still a proof-of-concept rather than a superior learning tool because networks with fixed weights could achieve similar performance. Attempts to evolve plastic neural controllers extended also to models of spiking neural networks (Maass and Bishop, 2001), which were tested in evolutionary learning experiments (Di Paolo, 2003; Federici, 2005). Spiking neural networks, however, require a more computationally intensive simulation, making evolution less feasible.

In short, ER sparked the discussion around design principles for robots and controllers, proposing the idea that bio-inspired evolutionary design may provide innovative solutions to control problems. In this attempt, evolving learning neural systems played a central role. Examples of the strong focus on learning in ER projects are provided, e.g., in The Cyber Rodent Project (Doya and Uchibe, 2005) that investigated the evolution of learning by seeking to implement a number of features such as (1) evolution of neural controllers, (2) learning of foraging and mating behaviors, (3) evolution of learning architectures and meta-parameters, (4) simultaneous learning of multiple agents in a body, and (5) learning and evolution in a self-sustained colony.

Experiments in ER in the 1990s and early 2000s revealed

the extent, complexity, and multitude of ideas behind the evolutionary design of learning robots. The challenges of evolving robotic hardware and control are combined with the challenges of setting the correct levels of abstraction for all parts of the solution, of evolving learning mechanisms, and of devising a conducive evolutionary and learning environment. For these reasons, the advances in ER led to a shift in focus more on specific problems. Among those are the interaction of evolution and learning, the evolution of neuromodulation, the evolution of learning architectures, and the evolution of indirectly encoded plasticity.

#### D. The interaction of evolution and learning

EPANNs are affected by the complex interaction of evolution and learning. Such interaction dynamics are present in any system that seeks the evolution of learning, and for this reason, it is worth presenting a few general concepts extending beyond EPANNs in the next paragraph.

Lifetime learning was first suggested to have an impact on accelerating evolution by Baldwin (1896). The studies in Smith (1986); Hinton and Nowlan (1987); Boers et al. (1995); Mayley (1996); Bullinaria (2001) demonstrated the Baldwin effect with computational simulations. For example, Mayley (1997) investigated the effects of learning on the rate of evolution, showing that a learning process during an individual lifetime, although costly, can improve the speed of evolution of the species. An important distinction can be drawn between (1) learning invariant environmental features and (2) learning variable environmental features. Hinton and Nowlan (1987) investigated the first case with a static environment and proved that learning can speed up the evolutionary acquisition of invariant knowledge by creating a gradient in the fitness landscape. In dynamic and uncertain environmental conditions, instead, knowledge cannot be acquired by the evolutionary process, as lifetime learning is required, and therefore the interaction of evolution and learning is less understood. A large body of work is dedicated to study the interaction and dynamics of all three POE dimensions, e.g. Paenke (2008); Paenke et al. (2009). Paenke et al. (2007) investigated the influence of plasticity and learning on evolution under directional selection; Mery and Burns (2010) demonstrate that behavioral plasticity affects evolution.

In EPANNs, similarly complex interaction dynamics between evolution and learning were also observed. One study (Soltoggio et al., 2007) suggests that evolution discovers learning by discrete stepping stones: when evolution casually discovers a weak mechanism of learning, it is sufficient to create an evolutionary advantage, so the neural mechanism is subsequently evolved quickly to perfect learning: Fig. 3 shows the best fitness in a foraging task when the agent suddenly evolves a learning strategy. When an environment changes over time, the frequency of those changes plays a role. With timescales comparable to a lifetime, evolution may lead to phenotypic plasticity, which is the capacity for a genotype to express different phenotypes in response to different environmental conditions (Lalejini and Ofria, 2016). The frequency of environmental changes was observed experimentally in




Figure 3: Discovery of a learning strategy during evolution. After approximately 1,000 generation, a reward signal is used by one network to learn where is the changing location of a stochastic reward. Figure adapted from Soltoggio et al. (2007).

plastic neural networks to affect the evolution of learning (Ellefsen, 2014), revealing a complex relationship between environmental variability and evolved learning.

A focus on the deceptiveness of evolving to learn is presented in Risi et al. (2010); Lehman and Miikkulainen (2014). In Risi et al. (2009, 2010), EPANN-controlled simulated robots, evolved in a discrete T-Maze domain, reveal that the evolutionary stepping stones towards adaptive behaviors are often not rewarded by objective-based performance measures. In other words, the necessary evolutionary stepping stones towards achieving adaptive behavior receive a lower fitness score than more brittle solutions with low learning but effective behaviors. A solution to this problem (Risi et al., 2010, 2009) was devised in Lehman and Miikkulainen (2014), in which *novelty search* (Lehman and Stanley, 2011) was adopted as a substitute to performance in the fitness objective with the aim of finding novel behaviors. Novelty search was observed to perform significantly better in the T-Maze domain.

Interestingly, by rewarding novel behaviors, novelty search validates the importance of exploration or curiosity, previously proposed in Schmidhuber (1991, 2006), also from an evolutionary viewpoint. With the aim of validating the same hypothesis, Soltoggio and Jones (2009) devised a simple EPANN experiment in which exploring was more advantageous than exploiting when reward information was not used, e.g., before learning evolves; to do this, the reward at a particular location depleted itself if continuously visited, so that changing location at random in a T-maze became beneficial. Interestingly, evolution discovered first exploratory behavior, even if not linked to learning, which in turn lead to an earlier evolution of reward-based learning. In that experiment, a stepping stone to evolve reward-based learning was to evolved non-reward based exploration.

The problem of deception in the evolution of learning can be seen as becoming trapped in a local optimum from which escaping requires the evolution of complex strategies. Such a situation emerges when an agent evolves a non-learning strategy that consistently and reliably delivers an average reward, as opposed to an agent that attempts learning, and by doing so risks collecting less reward during the learning






Figure 4: A modulatory neuron gates plasticity of the synapses that connect to the postsynaptic neuron. The learning is local, but a learning signal can be created by part of the network and used to regulate learning elsewhere.

process. The seminal work in Bullinaria (2003, 2007a, 2009b) proposes the more general hypothesis that learning requires the evolution of long periods of parental protection and late onset of maturity. Similarly, Ellefsen (2013b,a) investigates sensitive and critical periods of learning in evolved neural networks. This fascinating hypothesis has wider implications for experiments with EPANNs, and more generally for machine learning and AI. It is therefore foreseeable that future EPANNs will have a protected childhood during which parental guidance may be provided (Clutton-Brock, 1991; Klug and Bonsall, 2010).

#### E. The evolution of neuromodulation

Growing neuroscientific evidence on the role of neuromodulation (previously outlined in Section II-B) inspired the design of experiments with neuromodulated plasticity to evolve control behavior and learning strategies. A general form of neuromodulated plasticity can be expressed as

$$\Delta w = m \cdot f(x, y, w, \theta) \quad , \quad (3)$$

where  $m$  is a modulatory signal used as a multiplication factor to gate plasticity. Graphically, modulation can be represented as a different type of signal affecting plasticity at the synaptic connections of the afferent neuron (Fig. 4).

Evolutionary search was used to find the parameters of a neuromodulated Hebbian learning rule in a reward-based armed-bandit problem in Niv et al. (2002). The same problem was used later in Soltoggio et al. (2007) to evolve arbitrary learning architectures with a bio-inspired gene representation method called Analog Genetic Encoding (AGE) (Mattiussi and Floreano, 2007). Evolution was used to search modulatory topologies and parameters of a particular form of Eq. 3:

$$\Delta w = m \cdot (Axy + Bx + Cy + D) \quad , \quad (4)$$

where the parameters  $A$  to  $D$  determine the influence of four factors in the rule: a multiplicative Hebbian term  $A$ , a presynaptic term  $B$ , a postsynaptic term  $C$ , and pure modulatory, or heterosynaptic, term  $D$ .  $D$  allows the modulatory term alone to increase or decrease the afferent weights of the target neuron regardless of pre or postsynaptic activities. At the same time, Kondo (2007) proposed an evolutionary design and behavior analysis of neuromodulatory neural networks for mobile robot control, validating the potential of the method. Soltoggio et al. (2008) tested the hypothesis of whether modulatory dynamics held an evolutionary advantage. In their algorithm, modulatory

neurons are freely inserted or deleted by mutations, effectively allowing the system to autonomously select the computational components that give an evolutionary advantage. The evolving networks demonstrated that modulatory neurons were inserted into control networks by evolution with functional learning properties, providing an evolutionary advantage and better learning for those networks. In another study, Soltoggio (2008c) suggested that evolved modulatory neurons may be essential to separate the learning circuitry from the input-output controller, effectively contributing to shortening the input-output pathways and speeding up decision processes. The fundamental question of what learning rules and dynamics are required to solve which problems remains mostly unanswered (Soltoggio, 2008b), mainly because learning dynamics are affected by tight coupling between rules and architectures in a search space with many equivalent but different control structures. Fig.4 also suggests that modulatory networks require evolution to find two essential topological structures: what signals or combination of signals triggers modulation, and what neurons are to be targeted by modulatory signals.

A number of further studies on the evolution of neuromodulatory dynamics confirmed the evolutionary advantages in learning scenarios (Soltoggio, 2008a). Silva et al. (2012a) used simulations of 2-wheel robots performing a dynamic concurrent foraging task, in which scattered food items periodically change their nutritive value or become poisonous, similarly to the setup in Soltoggio and Stanley (2012). The results showed that when neuromodulation is enabled, learning is evolved faster than when neuromodulation is not enabled, also with multi-robot distributed systems (Silva et al., 2012b). Nogueira et al. (2013) also reported evolutionary advantages in foraging behavior of an autonomous virtual robot when equipped with neuromodulated plasticity. Harrington et al. (2013) demonstrated how evolved neuromodulation applied to a gene regulatory network consistently generalizes better than agents trained with fixed parameter settings. Surprisingly, Arnold et al. (2013b) showed that neuromodulatory architectures provide an evolutionary advantage also in reinforcement-free environments. The evolution of social representations in neural networks was shown to be facilitated by neuromodulatory dynamics in Arnold et al. (2013a). An artificial life simulation environment called Polyword (Yoder and Yaeger, 2014) helped to assess the advantage of neuromodulated plasticity in various scenarios. The authors found that neuromodulation may be able to enhance or diminish foraging performance in a competitive, dynamic environment.

Neuromodulation was evolved in Ellefsen et al. (2015) in combination with modularity to address the problem of catastrophic forgetting. In Gustafsson (2016), networks evolved with AGE (Mattiussi and Floreano, 2007) to play a game were shown to perform better with the addition of neuromodulation. In Norouzzadeh and Clune (2016), the authors demonstrated that neuromodulation produces forward models that can adapt to changes significantly better than the controls. They verified that evolution exploited variable learning rates to perform adaptation when needed.

The evidence in these studies suggests that neuromodulation is a key ingredient to facilitate the evolution of learning in

EPANN. They also indirectly suggest that neural systems with more than one type of signal (i.e. activation as well as other modulatory signals) are beneficial in the neuroevolution of learning.

#### F. Evolving learning architectures

Just as neuromodulation is one particular plasticity mechanism that enables effective reward learning, other principles may benefit particular types of learning behavior, as outlined in the review by (Floresano et al., 2008) and in a recent book by Downing (2015). Evolving architectures might be as simple as optimizing the number of hidden neurons, or the number of layers in a network, but in general, it refers to the discovery of neural pathways, learning rules, and other mechanisms whose synergetic matching enables particular learning dynamics. It is also important to note that EPANN do not have the neurophysiological constraints of biological systems, e.g., short neural connections, sparsity, brain size limits, etc., which impacts the evolution of particular architectures such as modularity (Bullinaria, 2007b, 2009a).

One seminal early study by Happel and Murre (1994) proposed the evolutionary design of modular neural networks in which modules could perform unsupervised learning, and the intermodule connectivity was shaped by Hebbian rules. In Khan et al. (2008), the authors propose an evolutionary developmental system. The resulting neural network is an example of an architecture that adapts with learning: the network has a dynamic morphology in which neurons can be inserted or deleted, and synaptic connections form and change in response to stimuli. Downing (2007) looks at different computational models of neurogenesis to propose an evolutionary developmental system, focusing in particular on abstraction levels and principles such as Neural Darwinism (Edelman and Tononi, 2000). In Khan et al. (2011); Khan and Miller (2014), the authors introduce a large number of bio-inspired mechanisms to evolve networks with rich learning dynamics. The idea is to use evolution to design a network that is capable of advanced plasticity such as dendrite branch and axon growth and shrinkage, neuron insertion and destruction, and many others.

In short, learning mechanisms and neural architectures appear strongly interdependent: EPANNs offer a tool to study the evolution of both simultaneously. This dependency is currently of particular interest also outside the field of EPANNs: in deep networks, learning rules and architectures need to be manually and carefully tuned to achieve good performance.

#### G. Indirectly Encoded Plasticity

Just as connection weights in EPANN may not be individually encoded in the genome, also plasticity may be coded in the genome to be expressed indirectly in the phenotype. Using a T-Maze domain as learning task, Risi and Stanley (2010) showed that HyperNEAT, which usually implements a compact encoding of weight patterns for large-scale ANNs (Fig. 5a), can also encode patterns of local learning rules. The approach, called *adaptive HyperNEAT*, can encode arbitrary learning rules for each connection in an evolving

ANN based on a function of the ANN's geometry (Fig. 5b). Further flexibility was added in Risi and Stanley (2012) to simultaneously encode the density and placement of nodes in substrate space. The approach, called *adaptive evolvable-substrate HyperNEAT*, makes it possible to indirectly encode plastic ANNs with thousands of connections that exhibit regularities and repeating motifs. Adaptive ES-HyperNEAT allows each individual synaptic connection, rather than neuron, to be standard or modulatory, thus introducing further design flexibility.

More recently, Risi and Stanley (2014) showed how adaptive HyperNEAT can be seeded to produce a specific lateral connectivity pattern, thereby allowing the weights to self-organize to form a genuine topographic map of the input space. This advance is important because it shows how evolution can be seeded with specific plasticity mechanisms that could facilitate the evolution of higher cognitive abilities.

The effect of encoding plasticity rules as patterns on the learning and on the evolutionary process was considered by Tonelli and Mouret (2013). Using an operant conditioning task, i.e., learning by reward, the authors showed that indirect encodings that produce more regular neural structures also improve the general EPANN learning abilities when compared to direct encodings.

### IV. OPEN CHALLENGES AND NEW DIRECTIONS

The body of work reviewed in the previous section contributed from different angles to produce the idea of EPANNs. The principles at play draw from a large variety of bio-inspired concepts and technological advances. Thus, new advances in AI, machine learning and neural networks are currently setting the groundwork for new directions for EPANNs. This section presents promising research themes that have recently emerged, and have the potential to extend and radically change the field of EPANN.

#### A. Levels of abstraction and representations for neural adaptivity

Choosing the right level of abstraction and the right representation (Bengio et al., 2013) are themes at the heart of many problems in AI. Because EPANN draw inspiration from many different research areas (see Fig. 1), the choices of abstraction level and representation are multiple.

Levels of abstraction are influenced by two main factors: low levels require increased computational costs, while high levels require an intuition of the essential dynamics that are necessary in the model. In the first case, the increase of computational power with modern hardware provides increasingly more resources to model large-scale neural systems at low levels, and to operate within complex environments. However, low level models are appropriate to model low level dynamics, but might not appropriately answer questions on the high level mechanisms on which evolution and learning operate.

One case in which the abstraction level is still clearly an open question is that of two fields simulating large neural networks: deep learning and bio-physical neural models. The progress in learning performance with deep networks

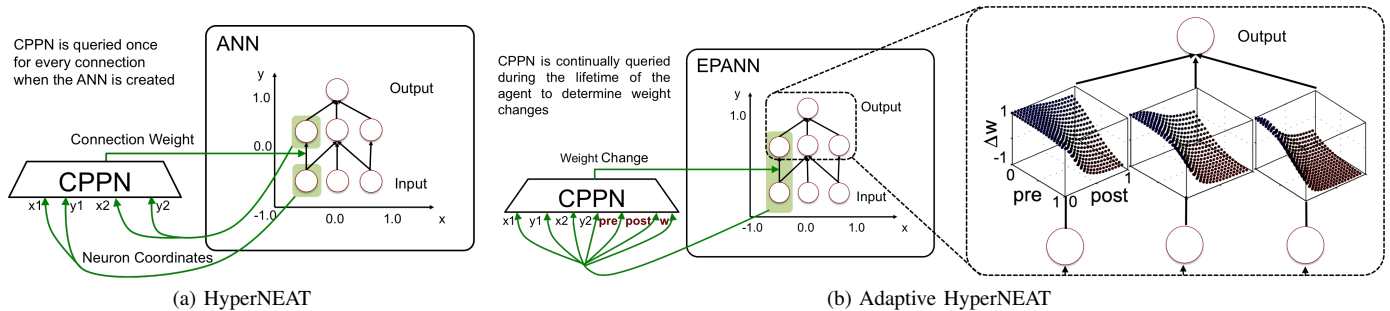


Figure 5: Example of an indirect mapping of plasticity rules from a compact genotype to a larger phenotype. (a) ANN nodes in HyperNEAT are situated in physical space by assigning them specific coordinates. The connections between nodes are determined by an evolved Compositional Patterns Producing Network (CPPN; Stanley (2007)), which takes as inputs the coordinates of two ANN neurons and returns the weight between them. In the normal HyperNEAT approach (a), the CPPN is queried once for all potential ANN connections when the agent is born. On the other hand, in adaptive HyperNEAT (b), the CPPN is continually queried during the lifetime of the agent to determine individual connection weight changes based on the location of neurons and additionally the activity of the presynaptic and postsynaptic neuron, and current connection weight. Adaptive HyperNEAT is able to indirectly encode a pattern of nonlinear learning rules for each connection in the ANN (right).

(Krizhevsky et al., 2012; LeCun et al., 2015; Schmidhuber, 2015), based mainly on rate-based neurons and low precision in the computation appears to contrast with the low-level simulations of accurate bio-physical models such as the Blue Brain Project (Markram, 2006), intended to reveal the computational and learning properties of a cortical column. Although the two approaches have different aims, they both represent abstractions of neural computation. Research in EPANN is best placed to address the problem of levels of abstractions because it can autonomously discover the most appropriate abstraction levels to evolve neural adaptation and intelligence.

In the context of genotype-to-phenotype mappings, Compositional Patterns Producing Networks (CPPNs) (Stanley, 2007), and also the previous work of Sims (1991), demonstrated that structured phenotypes can be generated through a function mapping without going through the dynamic developmental process typical of multicellular organisms. Relatedly, Hornby et al. (2002) showed that the different phenotypical representations led to considerably different results in the evolution of furniture. Miller (2014) discussed explicitly the effect of abstraction levels for evolved developmental learning networks, in particular in relation to two approaches that model development at the neuron level or at the network level.

This body of work suggests that finding appropriate representations, just as it was fundamental in the advances in deep learning to represent input spaces and hierarchical features (Bengio et al., 2013; Oquab et al., 2014), can also extend to representations of internal models, learning mechanisms, and genetic encodings, affecting the algorithms' capabilities of evolving learning abilities.

### B. Evolving general learning

One challenge in the evolution of learning is that evolved learning may simply result in the switch among a finite set of evolved behaviors, e.g. turning left or right in a T-Maze in a finite sequence, which is all that evolving solutions encounter during lifetime. A challenge for EPANNs is to acquire general

learning abilities in which the network is capable of learning problems not encountered during evolution. Mouret and Tonelli (2014) propose the distinction between the evolution of behavioral switches and the evolution of synaptic general learning abilities, and suggest conditions that favor these types of learning. General learning can be intuitively understood as the capability to learn any association among input, internal, and output patterns, both in the spatial and temporal dimensions, regardless of the complexity of the problem. Such objective clearly poses practical and philosophical challenges. Although humans are considered better at general learning than machines, human learning skills are specific and not unlimited (Ormrod and Davis, 2004). Nevertheless, moving from behavior switches to more general learning is a desirable feature for EPANN. Encouraging the emergence of general learners may likely involve (1) an increased computational cost for testing in rich environments, which include a large variety of uncertain and stochastic scenarios with problems of various complexity, and (2) an increased search space to explore the evolution of complex strategies and avoid deception.

One promising direction is to evolve the ability to learn algorithms and programs instead of simple associations. This approach has recently shown promise in the context of deep learning; methods such as the Neural Turing Machine (Graves et al., 2014) can learn simple algorithms such as sort, copy, etc. and approaches such as the Neural Programmer-Interface (Reed and de Freitas, 2015) can learn to represent and execute programs.

### C. Incremental and social learning

An important open challenge for machine learning in general is the creation of neural systems that can continuously integrate new knowledge and skills without forgetting what they previously learned. A promising deep learning approach is progressive neural networks (Rusu et al., 2016), in which a new network is created for each new task, and lateral connections between networks allow the system to leverage

previously learned features. Weight consolidation, formalized in a rule called Hypothesis Testing Plasticity, was shown in Soltoggio (2015) to permit learning of multiple tasks without the catastrophic forgetting typical of neural networks. The method suggested the consolidation of weights when associations became established after a period of learning. A similar idea was later named *elastic weight consolidation* in Kirkpatrick et al. (2017). Plasticity rules that implement weight consolidation (Soltoggio, 2015; Kirkpatrick et al., 2017) have proven to enhance networks to learn new tasks without suffering from catastrophic forgetting. Such approaches may become standard practice in EPANNs.

Encouraging modularity (Ellefsen et al., 2015; Durr et al., 2010) or augmenting evolving networks with a dedicated external memory component (Lüders et al., 2016) have been proposed recently. The scalability of these approaches and their impact on the evolutionary path will be a further topic of investigation. However, an evolutionary advantage is likely to emerge for networks that can elaborate on previously learned sub-skills during their lifetime to learn more complex tasks.

One interesting case in which incremental learning may play a role is social learning (Best, 1999). EPANNs may learn both from the environment and from other individuals, from scratch or incrementally (Offerman and Sonnemans, 1998). When social, learning may involve imitation, language or communication, or other social behaviors. Bullinaria (2017) proposes an EPANN framework to simulate the evolution of culture and social learning. It is reasonable to assume that future AI learning systems, whether based on EPANNs or not, will acquire knowledge through different modalities, involving direct experience with the environment, but also social interaction, and possibly with complex incremental learning phases.

#### D. Fast learning through generalization

Animal learning does not always require myriads of trials. Humans can very quickly generalize from only a few given examples, possibly leveraging previous experiences and a long learning process during infancy. This type of learning is currently missing in EPANNs. Inspiration for new approaches could come from complementary learning systems (Kumaran et al., 2016) that humans seem to possess, which include fast and slow learning components. Additionally, approaches such as probabilistic program induction seem to be able to learn concepts in one-shot at a human-level in some tasks (Lake et al., 2015). Fast learning is likely to derive not just from trial-and-error, but also from mental models that can be applied to diverse problems, similarly to transfer learning (Pan and Yang, 2010). Reusable mental models, once learned, will allow agents to make predictions and plan in new and uncertain scenarios with similarities to previously learned ones. If EPANNs can discover neural structures or learning rules that allow for generalization, an evolutionary advantage of such a discovery will lead to its full emergence and further optimization of such a property.

#### E. The loci of memory

The consequence of learning is memory, both explicit and implicit (Anderson, 2013), and its consolidation (Dudai, 2012). EPANNs may reach solutions in which memory evolved in different fashions, e.g., preserved as self-sustained neural activity, encoded by connection weights modified by plasticity rules, stored with an external memory (e.g. Neural Turing Machine), or a combination of these approaches. Neuromodulation and weight consolidation target areas of the network where information is stored (Soltoggio et al., 2008; Soltoggio, 2015). Recurrent neural architectures based on *long short-term memory* (LSTM) allow very complex tasks to be solved through gradient descent training (Greff et al., 2015; Hochreiter and Schmidhuber, 1997) and have recently shown promise when combined with evolution (Rawal and Miikkulainen, 2016). A different approach was introduced by Graves et al. (2014), called *Neural Turing Machine* (NTM), in which networks are augmented with an external memory that allows long-term memory storage. NTMs have shown promise when trained through evolution (Greve et al., 2016) or gradient descent (Graves et al., 2014, 2016). However, imposing such an explicit representation of memories may not lead to evolvable networks when problems require distributed, possibly implicit memory representations. Research in this area will reveal which computational systems are more evolvable and how memories will self organize and form in EPANN.

#### F. EPANN and deep learning

Deep learning has shown remarkable results in a variety of different fields (Krizhevsky et al., 2012; Schmidhuber, 2015; LeCun et al., 2015). However, the model structures of these networks are mostly hand-designed, include a large number of parameters, and require extensive experiments to discover optimal configurations. Advances in deep networks have come from manually testing various design aspects such as: different transfer functions, e.g. the ReLU (Krizhevsky et al., 2012); generalization mechanisms, e.g. the dropout (Srivastava et al., 2014); networks with different depth, e.g. (Simonyan and Zisserman, 2014), and others (Schmidhuber, 2015). With increased computational resources, it is now possible to search those design aspects with evolution, effectively in what can be called deep EPANNs. It is worth noting that the evolution of static weights for deep networks are not EPANNs. In contrast, certain types of EPANNs that evolve architectures with arbitrary depth, e.g. (Risi and Stanley, 2012), are effectively deep EPANNs.

Koutník et al. (2014) used evolution to design a controller that combined evolved recurrent neural networks, for the control part, and a deep max-pooling convolutional neural network to reduce the input dimensionality. The study does not use evolution on the deep preprocessing networks itself, but in combination with it, which nevertheless demonstrates the evolutionary design of a deep neural controller. Young et al. (2015) used an evolutionary algorithm to optimize two parameters of a deep network: the size (range [1,8]) and the number (range [16,126]) of the filters in a convolutional neural network. The study showed that the optimized parameters

can vary considerably from the standard best-practice values. An established evolutionary computation technique, the Covariance Matrix Adaptation Evolution Strategy (CMA-ES), was used in Loshchilov and Hutter (2016) to optimize the parameters of a deep network to learn to classify the MNIST dataset. The authors reported performance close to the state-of-the-art using 30 GPU devices. Fernando et al. (2016) was able to rediscover convolutional networks by means of evolution of Differentiable Pattern Producing Networks (Stanley, 2007).

Recently, Real et al. (2017) and Miikkulainen et al. (2017) showed that evolutionary search can be used to determine the topology, hyperparameters and building blocks of deep networks trained through gradient descent. The performance were shown to rival those of hand-designed architectures in the CIFAR-10 classification task and a language modeling task (Miikkulainen et al., 2017), while Real et al. (2017) also tested the method on the larger CIFAR-100 dataset. A different approach was presented by Fernando et al. (2017), in which evolution determines a subset of pathways through a network that are trained through backpropagation, allowing the same network to learn a variety of different tasks. While these results were only possible with significant computational resources, they demonstrate the potential of combining evolution and deep learning approaches.

The flexibility of evolutionary search, which can in some cases work better than gradient descent in search spaces with local optima, allows the design of search processes that not only tune certain parts of the networks, but also combine different mechanisms in an innovative and counter-intuitive fashion. An important point is that EPANNs do not necessarily implement error-based rules like backpropagation. Thus, the network need not be differentiable, as it would need to be in gradient descent based approaches, thereby allowing a greater variety in architectures (Greve et al., 2016). Finally, the deception in the evolution of learning suggests that gradient descent approaches can encounter limitations when searching for complex learning strategies, leaving EPANN as a compelling alternative to discover more advanced cognitive learning models.

#### *G. GPU implementations and neuromorphic hardware*

The progress of EPANNs will crucially depend on implementations that take advantage of the increased computational power of parallel computation with GPUs and neuromorphic hardware (Jo et al., 2010; Monroe, 2014). Deep learning greatly benefited from GPU-accelerated machine learning but also standardized tools (e.g. Torch, Tensorflow, Theano, etc.) that made it easy for anybody to download, experiment, and extend promising deep learning models.

Howard et al. (2011, 2012, 2014) devised experiments to evolve plastic spiking networks implemented as memristors for simulated robotic navigation tasks. Memristive plasticity was observed consistently to enable higher performance than constant-weighted connections in both static and dynamic reward scenarios. This body of work consolidates the idea that EPANNs have considerable advantages over non-learning networks.

In the context of newly emerging technologies, it is worth noting that, just as GPUs were not developed initially for deep learning, so novel neural computation tools and hardware systems, not developed for EPANNs, can now be exploited to setup the articulated processes and computationally expensive experiments required by EPANN.

#### *H. Measuring progress*

Much of the recent progress in neural learning comes from having standardized benchmarks and training sets that algorithms can be evaluated on. However, the measurement rarely considers the learning process itself, but rather the performance when the learning has ended. The development of criteria to measure learning skills (Miconi, 2008), specifically designed to measure the agent's learning capabilities, and not only a task-specific performance, is necessary to appreciate the scientific significance of, and progress in, EPANNs. A number of platforms for testing AI controllers are becoming available, e.g., the Atari or General Video Game Playing Benchmark (GVGAI, 2017), or the OpenAI Universe, "a software platform for measuring and training an AI's general intelligence across the world's supply of games, websites and other applications" (OpenAI, 2017). For EPANN using such benchmarks, the goal is not to be good at playing one or more games, but rather to test the capability to evolve the learning required for such control problems. Measuring such skills involves not only looking at the final performance, but also at the dynamics of learning (where fast learning is not always better), at the robustness, generalization, initial conditions, type of datasets or richness of the environments, size of the network, and computational resources used. Complex benchmarks could allow testing increasingly complex problems or scenarios, requiring the agent to build upon previously acquired skills.

### V. CONCLUSION

This paper first described the broad inspiration and aspirations of evolved plastic neural networks (EPANNs), which are shown to draw from large, diverse, and interdisciplinary areas. The inspirations also reveal the ambitious and long-term research objectives of EPANNs with important implications for Artificial Intelligence and biology.

EPANNs saw considerable progress in the last two decades, which led to the discovery of a number of further research themes, such as the design of evolutionary algorithms to promote the evolution of learning, a better understanding of the complex interaction dynamics between evolution and learning, the advantages of multi-signal networks such as modulatory networks, and the attempt to evolve architectures and appropriate representations of learning mechanisms.

Finally, the paper outlined that recent scientific and technical progress may enable a step change in the potential of EPANNs. In light of current advances, a pivotal moment for EPANNs may be at hand as the next AI tool to find flexible algorithms, capable of discovering principles for general adaptation and intelligent systems.

## REFERENCES

- Abbott, L. F. (1990). Modulation of Function and Gated Learning in a Network Memory. *Proceedings of the National Academy of Science of the United States of America*, 87(23):9241–9245.
- Abraham, A. (2004). Meta learning evolutionary artificial neural networks. *Neurocomputing*, 56:1–38.
- Abraham, W. C. and Robins, A. (2005). Memory retention—the synaptic stability versus plasticity dilemma. *Trends in neurosciences*, 28(2):73–78.
- Alexander, W. H. and Sporns, O. (2002). An Embodied Model of Learning, Plasticity, and Reward. *Adaptive Behavior*, 10(3-4):143–159.
- Allis, L. V. et al. (1994). *Searching for solutions in games and artificial intelligence*. Ponsen & Looijen.
- Alpaydin, E. (2014). *Introduction to machine learning*. MIT press.
- Anderson, J. R. (2013). *The architecture of cognition*. Psychology Press.
- Angeline, P. J., Saunders, G. M., and Pollack, J. B. (1994). An evolutionary algorithm that constructs recurrent neural networks. *IEEE transactions on Neural Networks*, 5(1):54–65.
- Arnold, S., Suzuki, R., and Arita, T. (2013a). Evolution of social representation in neural networks. In *Advances in Artificial Life, ECAL 2013: Proceedings of the Twelfth European Conference on the Synthesis and Simulation of Living Systems*. MIT Press, Cambridge, MA, pages 425–430.
- Arnold, S., Suzuki, R., and Arita, T. (2013b). Selection for reinforcement-free learning ability as an organizing factor in the evolution of cognition. *Advances in Artificial Intelligence*, 2013:8.
- Auerbach, J., Fernando, C., and Floreano, D. (2014). Online extreme evolutionary learning machines. In *Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*.
- Back, T., Fogel, D. B., and Michalewicz, Z. (1997). *Handbook of evolutionary computation*. IOP Publishing Ltd.
- Bäck, T. and Schwefel, H.-P. (1993). An overview of evolutionary algorithms for parameter optimization. *Evolutionary computation*, 1(1):1–23.
- Bahroun, Y. and Soltoggio, A. (2017). Online Representation Learning with Multi-layer Hebbian Networks for Image Classification Tasks.
- Bailey, C. H., Giustetto, M., Huang, Y.-Y., Hawkins, R. D., and Kandel, E. R. (2000). Is heterosynaptic modulation essential for stabilizing Hebbian plasticity and memory? *Nature Reviews Neuroscience*, 1(1):11–20.
- Baldwin, J. M. (1896). A new factor in evolution (continued). *American naturalist*, pages 536–553.
- Baxter, D. A., Canavier, C. C., Clark, J. W., and Byrne, J. H. (1999). Computational Model of the Serotonergic Modulation of Sensory Neurons in Aplysia. *Journal of Neurophysiology*, 82:1914–2935.
- Baxter, J. (1992). The Evolution of Learning Algorithms for Artificial Neural Networks. *Complex Systems*, D. Green and T. Bossomaier, Eds. Amsterdam, The Netherlands: IOS, pages 313–326.
- Beer, R. D. and Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, 1(1):91–122.
- Bengio, S., Bengio, Y., Cloutier, J., and Gecsei, J. (1992). On the optimization of a synaptic learning rule. In *Preprints Conf. Optimality in Artificial and Biological Neural Networks*, Univ. of Texas, Dallas, Feb 6-8, 1992.
- Bengio, Y., Bengio, S., and Cloutier, J. (1990). Learning a synaptic learning rule. Technical report, Université de Montréal, Département d’informatique et de recherche opérationnelle.
- Bengio, Y., Courville, A., and Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828.
- Best, M. L. (1999). How culture can guide evolution: An inquiry into gene/meme enhancement and opposition. *Adaptive Behavior*, 7(3-4):289–306.
- Bienenstock, L. E., Cooper, L. N., and Munro, P. W. (1982). Theory for the development of neuron selectivity: orientation specificity and binocular interaction in visual cortex. *The Journal of Neuroscience*, 2(1):32–48.
- Birmingham, J. T. (2001). Increasing Sensor Flexibility Through Neuromodulation. *Biological Bulletin*, 200:206–210.
- Blynel, J. and Floreano, D. (2002). Levels of dynamics and adaptive behavior in evolutionary neural controllers. In *Proceedings of the seventh international conference on simulation of adaptive behavior on From animals to animats*, pages 272–281. MIT Press Cambridge, MA, USA.
- Blynel, J. and Floreano, D. (2003). Exploring the T-Maze: Evolving Learning-Like Robot Behaviors Using CTRNNs. In *EvoWorkshops*, pages 593–604.
- Boers, E. J., Borst, M. V., and Sprinkhuizen-Kuyper, I. G. (1995). Evolving neural networks using the “baldwin effect”. In *Artificial Neural Nets and Genetic Algorithms*, pages 333–336. Springer.
- Bourlard, H. and Kamp, Y. (1988). Auto-association by multilayer perceptrons and singular value decomposition. *Biological cybernetics*, 59(4-5):291–294.
- Brown, T. H., Kairiss, E. W., and Keenan, C. L. (1990). Hebbian Synapse: Biophysical Mechanisms and Algorithms. *Annual Review of Neuroscience*, 13:475–511.
- Bullinaria, J. A. (2001). Exploring the baldwin effect in evolving adaptable control systems. In *Connectionist models of learning, development and evolution*, pages 231–242. Springer.
- Bullinaria, J. A. (2003). From biological models to the evolution of robot control systems. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, pages 2145–2164.
- Bullinaria, J. A. (2007a). The effect of learning on life history evolution. In *Proceedings of the 9th annual conference on Genetic and evolutionary computation*, pages 222–229. ACM.
- Bullinaria, J. A. (2007b). Understanding the emergence of



- modularity in neural systems. *Cognitive science*, 31(4):673–695.
- Bullinaria, J. A. (2009a). The importance of neurophysiological constraints for modelling the emergence of modularity. *Computational modelling in behavioural neuroscience: closing the gap between neurophysiology and behaviour*, 2:188–208.
- Bullinaria, J. A. (2009b). Lifetime learning as a factor in life history evolution. *Artificial Life*, 15(4):389–409.
- Bullinaria, J. A. (2017). Imitative and direct learning as interacting factors in life history evolution. *Artificial Life*, 23.
- Carew, T. J., Walters, E. T., and Kandel, E. R. (1981). Classical conditioning in a simple withdrawal reflex in *Aplysia californica*. *The Journal of Neuroscience*, 1(12):1426–1437.
- Carver, C. S. and Scheier, M. F. (2012). *Attention and self-regulation: A control-theory approach to human behavior*. Springer Science & Business Media.
- Cervier, D. (1993). AI: The Tumultuous Search for Artificial Intelligence.
- Chalmers, D. J. (1990). The evolution of learning: An experiment in genetic connectionism. In *Proceedings of the 1990 connectionist models summer school*, pages 81–90. San Mateo, CA.
- Chiel, H. J. and Beer, R. D. (1997). The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment. *Trends in Neurosciences*, 20(12):553–557.
- Clark, G. A. and Kandel, E. R. (1984). Branch-specific heterosynaptic facilitation in *Aplysia* siphon sensory cells. *PNAS*, 81(8):2577–2581.
- Cliff, D., Husbands, P., and Harvey, I. (1993). Explorations in evolutionary robotics. *Adaptive Behavior*, 2(1):73–110.
- Clutton-Brock, T. H. (1991). *The evolution of parental care*. Princeton University Press.
- Cooper, S. J. (2005). Donald O. Hebb’s synapse and learning rule: a history and commentary. *Neuroscience and Biobehavioral Reviews*, 28(8):851–874.
- Damasio, A. R. (1999). *The feeling of what happens: Body and emotion in the making of consciousness*. Houghton Mifflin Harcourt.
- Darwin, C. (1859). *On the origin of species by means of natural selection, or The preservation of favoured races in the struggle for life*. Murray, London.
- Dawkins, R. (2003). The evolution of evolvability. *On growth, form and computers*, pages 239–255.
- De Botton, A. (2016). Six areas that artificial emotional intelligence will revolutionise. <http://www.wired.co.uk/article/wired-world-2016-alain-de-botton-artificial-emotional-intelligence>. Accessed: 12/11/2016.
- Deary, I. J., Johnson, W., and Houlihan, L. M. (2009). Genetic foundations of human intelligence. *Human genetics*, 126(1):215–232.
- Deng, L., Hinton, G., and Kingsbury, B. (2013). New types of deep neural network learning for speech recognition and related applications: An overview. In *Acoustics, Speech and Signal Processing (ICASSP), 2013 IEEE International Conference on*, pages 8599–8603. IEEE.
- Di Paolo, E. A. (2003). Evolving spike-timing-dependent plasticity for single-trial learning in robots. *Philosophical Transactions of the Royal Society*, 361(1811):2299–2319.
- Doidge, N. (2007). *The brain that changes itself: Stories of personal triumph from the frontiers of brain science*. Penguin.
- Downing, K. L. (2007). Supplementing Evolutionary Development Systems with Abstract Models of Neurogenesis. In at al., D. T., editor, *Proceeding of the Genetic and Evolutionary Computation Conference (GECCO 2007)*, pages 990–998. ACM.
- Downing, K. L. (2015). *Intelligence Emerging: Adaptivity and Search in Evolving Neural Systems*. MIT Press.
- Doya, K. (2002). Metalearning and neuromodulation. *Neural Networks*, 15(4-6):495–506.
- Doya, K. and Uchibe, E. (2005). The Cyber Rodent Project: Exploration and Adaptive Mechanisms for Self-Preservation and Self-Reproduction. *Adaptive Behavior*, 13(2):149–160.
- Draganski, B. and May, A. (2008). Training-induced structural changes in the adult human brain. *Behavioural brain research*, 192(1):137–142.
- Dudai, Y. (2012). The restless engram: consolidations never end. *Annual review of neuroscience*, 35:227–247.
- Durr, P., Floreano, D., and Mattiussi, C. (2010). Genetic representation and evolvability of modular neural controllers. *IEEE Computational Intelligence Magazine*, 5(3):10–19.
- Edelman, G. M. and Tononi, G. (2000). *A universe of consciousness: How matter becomes imagination*. Basic books.
- Eiben, A. E. and Smith, J. (2015). From evolutionary computation to the evolution of things. *Nature*, 521(7553):476–482.
- Eliasmith, C., Stewart, T. C., Choo, X., Bekolay, T., DeWolf, T., Tang, Y., and Rasmussen, D. (2012). A large-scale model of the functioning brain. *science*, 338(6111):1202–1205.
- Ellefsen, K. O. (2013a). Different regulation mechanisms for evolved critical periods in learning. In *2013 IEEE Congress on Evolutionary Computation*, pages 1193–1200. IEEE.
- Ellefsen, K. O. (2013b). Evolved sensitive periods in learning. *Advances in Artificial Life, ECAL 2013: Proceedings of the Twelfth European Conference on the Synthesis and Simulation of Living Systems*, 12:409–416.
- Ellefsen, K. O. (2014). The evolution of learning under environmental variability. *Artificial Life 14: Proceedings of The Fourteenth International Conference on the Synthesis and Simulation of Living Systems*.
- Ellefsen, K. O., Mouret, J.-B., Clune, J., and Bongard, J. C. (2015). Neural modularity helps organisms evolve to learn new skills without forgetting old skills. *PLoS Comput Biol*, 11(4):e1004128.
- Federici, D. (2005). Evolving Developing Spiking Neural Networks. In *Proceedings of the IEEE Congress on Evolutionary Computation, CEC 2005*.
- Fellous, J.-M. and Linster, C. (1998). Computational Models of Neuromodulation. *Neural Computation*, 10:771–805.
- Fernando, C., Banarse, D., Blundell, C., Zwols, Y., Ha, D., Rusu, A. A., Pritzel, A., and Wierstra, D. (2017). Pathnet: Evolution channels gradient descent in super neural

- networks. *arXiv preprint arXiv:1701.08734*.
- Fernando, C., Banarse, D., Reynolds, M., Besse, F., Pfau, D., Jaderberg, M., Lanctot, M., and Wierstra, D. (2016). Convolution by evolution: Differentiable pattern producing networks. *arXiv preprint arXiv:1606.02580*.
- Finnie, P. S. and Nader, K. (2012). The role of metaplasticity mechanisms in regulating memory destabilization and reconsolidation. *Neuroscience & Biobehavioral Reviews*, 36(7):1667–1707.
- Floreano, D., Dürr, P., and Mattiussi, C. (2008). Neuroevolution: from architectures to learning. *Evolutionary Intelligence*, 1(1):47–62.
- Floreano, D. and Mattiussi, C. (2008). *Bio-inspired artificial intelligence: theories, methods, and technologies*. MIT press.
- Floreano, D. and Mondada, F. (1994). Automatic creation of an autonomous agent: genetic evolution of a neural-network driven robot. In *Proceedings of the third international conference on Simulation of adaptive behavior : from animals to animats 3*, pages 421–430. MIT Press, Cambridge, MA, USA.
- Floreano, D. and Mondada, F. (1996). Evolution of homing navigation in a real mobile robot. *IEEE Transactions on Systems, Man, and Cybernetics—Part B*, 26(3):396–407.
- Floreano, D., Mondada, F., Perez-Urbe, A., and Roggen, D. (2004). Evolution of embodied intelligence. In *Embodied artificial intelligence*, pages 293–311. Springer.
- Floreano, D. and Nolfi, S. (2004). *Evolutionary Robotics*. The MIT Press.
- Floreano, D. and Urzelai, J. (2001a). Evolution of plastic control networks. *Auton. Robots*, 11(3):311–317.
- Floreano, D. and Urzelai, J. (2001b). Neural morphogenesis, synaptic plasticity, and evolution. *Theory in Biosciences*, 240:225–240.
- Fogel, D. B. (2006). *Evolutionary computation: toward a new philosophy of machine intelligence*, volume 1. John Wiley & Sons.
- Fontanari, J. and Meir, R. (1991). Evolving a learning algorithm for the binary perceptron. *Network: Computation in Neural Systems*, 2(4):353–359.
- Fujii, A., Saito, N., Nakahira, K., and Ishiguro, A. (2002). Generation of an Adaptive Controller CPG for a Quadruped Robot with Neuromodulation Mechanism. In *Proceedings of the 2002 IEEE/RSJ International Conference On Intelligent Robots and Systems, EPFL, Lausanne, Switzerland*, pages 2619–2624.
- Gerstner, W. and Kistler, W. M. (2002a). Mathematical formulations of hebbian learning. *Biological cybernetics*, 87(5-6):404–415.
- Gerstner, W. and Kistler, W. M. (2002b). *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge university press.
- Goldberg, D. E. and Holland, J. H. (1988). Genetic algorithms and machine learning. *Machine learning*, 3(2):95–99.
- Graves, A., Wayne, G., and Danihelka, I. (2014). Neural turing machines. *arXiv preprint arXiv:1410.5401*.
- Graves, A., Wayne, G., Reynolds, M., Harley, T., Danihelka, I., Grabska-Barwińska, A., Colmenarejo, S. G., Grefenstette, E., Ramalho, T., Agapiou, J., et al. (2016). Hybrid computing using a neural network with dynamic external memory. *Nature*, 538(7626):471–476.
- Greff, K., Srivastava, R. K., Koutník, J., Steunebrink, B. R., and Schmidhuber, J. (2015). Lstm: A search space odyssey. *arXiv preprint arXiv:1503.04069*.
- Greve, R. B., Jacobsen, E. J., and Risi, S. (2016). Evolving neural turing machines for reward-based learning. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016, GECCO '16*, pages 117–124, New York, NY, USA. ACM.
- Grossberg, S. (2012). *Studies of mind and brain: Neural principles of learning, perception, development, cognition, and motor control*, volume 70. Springer Science & Business Media.
- Gu, Q. (2002). Neuromodulatory transmitter systems in the cortex and their role in cortical plasticity. *Neuroscience*, 111(4):815–853.
- Gustafsson, J. (2016). Evolving neuromodulatory topologies for plasticity in video game playing. Master’s thesis, Blekinge Institute of Technology, Faculty of Computing, Department of Software Engineering.
- GVGAI (2017). The general video game ai competition. Accessed: 21/01/2017.
- Happel, B. L. and Murre, J. M. (1994). Design and evolution of modular neural network architectures. *Neural networks*, 7(6):985–1004.
- Harrington, K. I., Awa, E., Cussat-Blanc, S., and Pollack, J. (2013). Robot coverage control by evolved neuromodulation. In *Neural Networks (IJCNN), The 2013 International Joint Conference on*, pages 1–8. IEEE.
- Harris-Warrick, R. M. and Marder, E. (1991). Modulation of neural networks for behavior. *Annual Review of Neuroscience*, 14:39–57.
- Harvey, I., Husbands, P., Cliff, D., Thompson, A., and Jakobi, N. (1997). Evolutionary robotics: the sussex approach. *Robotics and autonomous systems*, 20(2):205–224.
- Hasselmo, M. E. (1995). Neuromodulation and cortical function: modeling the physiological basis of behavior. *Behavioural Brain Research*, 67:1–27.
- Hasselmo, M. E. (1999). Neuromodulation: acetylcholine and memory consolidation. *Trends in Cognitive Sciences*, 3:351–359.
- Hasselmo, M. E. and Schnell, E. (1994). Laminar selectivity of the cholinergic suppression of synaptic transmission in rat hippocampal region CA1: computational modeling and brain slice physiology. *Journal of Neuroscience*, 14(6):3898–3914.
- Hawkins, J. and Blakeslee, S. (2007). *On intelligence*. Macmillan.
- Hebb, D. O. (1949). *The Organization of Behavior: A Neuropsychological Theory*. Wiley, New York.
- Hensch, T. K., Fagiolini, M., Mataga, N., Stryker, M. P., Baekkeskov, S., and Kash, S. F. (1998). Local gaba circuit control of experience-dependent plasticity in developing visual cortex. *Science*, 282(5393):1504–1508.
- Hinton, G. E. and Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, 1:495–502.

- Hinton, G. E. and Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science*, 313(5786):504–507.
- Hochreiter, S. and Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Holland, J. H. (1975). *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. U Michigan Press.
- Holland, J. H. and Reitman, J. S. (1977). Cognitive systems based on adaptive algorithms. *ACM SIGART Bulletin*, 63:49–49.
- Hopkins, W. D., Russell, J. L., and Schaeffer, J. (2014). Chimpanzee intelligence is heritable. *Current Biology*, 24(14):1649–1652.
- Hornby, G., Pollack, J., et al. (2002). Creating high-level components with a generative representation for body-brain evolution. *Artificial life*, 8(3):223–246.
- Howard, G., Bull, L., de Lacy Costello, B., Gale, E., and Adamatzky, A. (2014). Evolving spiking networks with variable resistive memories. *Evolutionary computation*, 22(1):79–103.
- Howard, G., Gale, E., Bull, L., de Lacy Costello, B., and Adamatzky, A. (2011). Towards evolving spiking networks with memristive synapses. In *Artificial Life (ALIFE), 2011 IEEE Symposium on*, pages 14–21. IEEE.
- Howard, G., Gale, E., Bull, L., de Lacy Costello, B., and Adamatzky, A. (2012). Evolution of plastic learning in spiking networks via memristive connections. *IEEE Transactions on Evolutionary Computation*, 16(5):711–729.
- Hull, C. L. (1943). *Principles of behavior*. New-York: Appleton Century.
- Husbands, P., Smith, T., Jakobi, N., and O’Shea, M. (1998). Better Living Through Chemistry: Evolving GasNet for Robots Control. *Connection Science*, 10:185–210.
- Izhikevich, E. M. (2007). Solving the distal reward problem through linkage of stdp and dopamine signaling. *Cerebral cortex*, 17(10):2443–2452.
- Jo, S. H., Chang, T., Ebong, I., Bhadviya, B. B., Mazumder, P., and Lu, W. (2010). Nanoscale memristor device as synapse in neuromorphic systems. *Nano letters*, 10(4):1297–1301.
- Kandel, E. R. (2007). *In Search of Memory - The Emergence of a New Science of Mind*. WW Norton & Company, New York.
- Kandel, E. R. and Tauc, L. (1965). Heterosynaptic facilitation in neurones of the abdominal ganglion of *Aplysia depilans*. *The Journal of Physiology*, 181:1–27.
- Katz, P. S. (1995). Intrinsic and extrinsic neuromodulation of motor circuits. *Current Opinion in Neurobiology*, 5:799–808.
- Katz, P. S. and Frost, W. N. (1996). Intrinsic neuromodulation: altering neuronal circuits from within. *Trends in Neurosciences*, 19:54–61.
- Katz, P. S., Getting, P. A., and Frost, W. N. (1994). Dynamic neuromodulation of synaptic strength intrinsic to a central pattern generator circuit. *Nature*, 367:729–731.
- Khan, G. M. and Miller, J. F. (2014). In search of intelligence: evolving a developmental neuron capable of learning. *Connection Science*, 26(4):297–333.
- Khan, G. M., Miller, J. F., and Halliday, D. M. (2008). Coevolution of neuro-developmental programs that play checkers. In Springer, editor, *Proceedings of 8th International Conference on Evolvable Systems (ICES): From Biology to Hardware*, volume 5216, pages 352–361.
- Khan, G. M., Miller, J. F., and Halliday, D. M. (2011). Evolution of cartesian genetic programs for development of learning neural architecture. *Evolutionary Computation*, 19(3):469–523.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., Clopath, C., Kumaran, D., and Hadsell, R. (2017). Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*.
- Kiyota, T. (2017). Neurogenesis and brain repair. In *Neuroimmune Pharmacology*, pages 575–597. Springer.
- Klug, H. and Bonsall, M. B. (2010). Life history and the evolution of parental care. *Evolution*, 64(3):823–835.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological cybernetics*, 43(1):59–69.
- Kohonen, T. (1990). The self-organizing map. *Proceedings of the IEEE*, 78(9):1464–1480.
- Kolb, B. (1989). Brain development, plasticity, and behavior. *American Psychologist*, 44(9):1203.
- Kolb, B. and Gibb, R. (2011). Brain plasticity and behaviour in the developing brain. *Journal of the Canadian Academy of Child & Adolescent Psychiatry*, 20(4).
- Kondo, T. (2007). Evolutionary design and behaviour analysis of neuromodulatory neural networks for mobile robots control. *Applied Soft Computing*, 7(1):189–202.
- Koutník, J., Schmidhuber, J., and Gomez, F. (2014). Evolving deep unsupervised convolutional networks for vision-based reinforcement learning. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*, pages 541–548. ACM.
- Koza, J. R. (1992). *Genetic programming: on the programming of computers by means of natural selection*, volume 1. MIT press.
- Krichmar, J. L. (2008). The neuromodulatory system: a framework for survival and adaptive behavior in a challenging world. *Adaptive Behavior*, 16(6):385–399.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105.
- Kumaran, D., Hassabis, D., and McClelland, J. L. (2016). What learning systems do intelligent agents need? complementary learning systems theory updated. *Trends in Cognitive Sciences*, 20(7):512–534.
- Kupfermann, I. (1987). Cellular neurobiology: Neuromodulation. *Science*, 236:863.
- KϕPPE, S. (1983). The psychology of the neuron: Freud, cajal and golgi. *Scandinavian journal of psychology*, 24(1):1–12.
- Lake, B. M., Salakhutdinov, R., and Tenenbaum, J. B. (2015). Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332–1338.
- Lalejini, A. and Ofria, C. (2016). The evolutionary origins of

- phenotypic plasticity. In *Proceedings of the Artificial Life Conference*.
- Lamprecht, R. and LeDoux, J. (2004). Structural plasticity and memory. *Nature Reviews Neuroscience*, 5(1):45–54.
- Langton, C. G. (1997). *Artificial life: An overview*. Mit Press.
- LeCun, Y., Bengio, Y., and Hinton, G. (2015). Deep learning. *Nature*, 521(7553):436–444.
- LeDoux, J. E. (2003). *Synaptic self: How our brains become who we are*. Penguin.
- Lee, C. M. and Narayanan, S. S. (2005). Toward detecting emotions in spoken dialogs. *IEEE transactions on speech and audio processing*, 13(2):293–303.
- Lehman, J. and Miikkulainen, R. (2014). Overcoming deception in evolution of cognitive behaviors. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2014)*, Vancouver, BC, Canada.
- Lehman, J. and Stanley, K. O. (2008). Exploiting open-endedness to solve problems through the search for novelty. In *ALIFE*, pages 329–336.
- Lehman, J. and Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223.
- Lehman, J. and Stanley, K. O. (2013). Evolvability is inevitable: increasing evolvability without the pressure to adapt. *PloS one*, 8(4):e62186.
- Levy, J. P. and Bairaktaris, D. (1995). Connectionist dual-weight architectures. *Language and Cognitive Processes*, 10(3-4):265–283.
- Lin, L.-J. (1993). *Reinforcement Learning for Robots Using Neural Networks*. PhD thesis, School of Computer Science, Carnegie Mellon University.
- Lisman, J. (1989). A mechanism for the Hebb and the anti-Hebb processes underlying learning and memory. *PNAS*, 86:9574–9578.
- Loshchilov, I. and Hutter, F. (2016). Cma-es for hyperparameter optimization of deep neural networks. *arXiv preprint arXiv:1604.07269*.
- Lüders, B., Schläger, M., and Risi, S. (2016). Continual learning through evolvable neural turing machines. In *NIPS 2016 Workshop on Continual Learning and Deep Networks (CLDL 2016)*.
- Maass, W. and Bishop, C. M. (2001). *Pulsed neural networks*. MIT press.
- Marder, E. (1996). Neural modulation: Following your own rhythm. *Current Biology*, 6(2):119–121.
- Marder, E. and Thirumalai, V. (2002). Cellular, synaptic and network effects of neuromodulation. *Neural Networks*, 15:479–493.
- Markram, H. (2006). The blue brain project. *Nature Reviews Neuroscience*, 7(2):153–160.
- Markram, H. (2012). The human brain project. *Scientific American*, 306(6):50–55.
- Markram, H., Lübke, J., Frotscher, M., and Sakmann, B. (1997). Regulation of synaptic efficacy by coincidence of postsynaptic aps and epsps. *Science*, 275(5297):213–215.
- Mattiussi, C. and Floreano, D. (2007). Analog genetic encoding for the evolution of circuits and networks. *IEEE Transactions on evolutionary computation*, 11(5):596–607.
- Mayley, G. (1996). Landscapes, learning costs, and genetic assimilation. *Evolutionary Computation*, 4(3):213–234.
- Mayley, G. (1997). Guiding or hiding: Explorations into the effects of learning on the rate of evolution. In *Proceedings of the Fourth European Conference on Artificial Life*, volume 97, pages 135–144.
- McQuesten, P. and Miikkulainen, R. (1997). Culling and teaching in neuro-evolution. In *ICGA*, pages 760–767.
- Mery, F. and Burns, J. G. (2010). Behavioural plasticity: an interaction between evolution and experience. *Evolutionary Ecology*, 24(3):571–583.
- Merzenich, M. M., Nelson, R. J., Stryker, M. P., Cynader, M. S., Schoppmann, A., and Zook, J. M. (1984). Somatosensory cortical map changes following digit amputation in adult monkeys. *Journal of comparative neurology*, 224(4):591–605.
- Michalewicz, Z. (1994). Gas: What are they? In *Genetic algorithms+ data structures= evolution programs*, pages 13–30. Springer.
- Michalski, R. S., Carbonell, J. G., and Mitchell, T. M. (2013). *Machine learning: An artificial intelligence approach*. Springer Science & Business Media.
- Miconi, T. (2008). *The road to everywhere: Evolution, complexity and progress in natural and artificial systems*. PhD thesis, University of Birmingham.
- Miikkulainen, R., Liang, J., Meyerson, E., Rawal, A., Fink, D., Francon, O., Raju, B., Navruzian, A., Duffy, N., and Hodjat, B. (2017). Evolving deep neural networks. *arXiv preprint arXiv:1703.00548*.
- Miller, J. F. (2014). Neuro-centric and holocentric approaches to the evolution of developmental neural networks. In *Growing Adaptive Machines*, pages 227–249. Springer.
- Miller, K. D. and Mackay, D. J. C. (1994). The Role of Constraints in Hebbian Learning. *Neural Computation*, 6:100–126.
- Millington, I. and Funge, J. (2016). *Artificial intelligence for games*. CRC Press.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533.
- Monroe, D. (2014). Neuromorphic computing gets ready for the (really) big time. *Communications of the ACM*, 57(6):13–15.
- Mouret, J.-B. and Tonelli, P. (2014). Artificial evolution of plastic neural networks: a few key concepts. In *Growing Adaptive Machines*, pages 251–261. Springer.
- Niv, Y., Joel, D., Meilijson, I., and Rupp, E. (2002). Evolution of Reinforcement Learning in Uncertain Environments: A Simple Explanation for Complex Foraging Behaviours. *Adaptive Behavior*, 10(1):5–24.
- Nogueira, B., Lenon, Y., Eduardo, C., Vidal, C. A., and Cavalcante Neto, J. B. (2013). Evolving plastic neuromodulated networks for behavior emergence of autonomous virtual characters. In *Advances in Artificial Life, ECAL 2013: Proceedings of the Twelfth European Conference on the Synthesis and Simulation of Living Systems*. MIT Press,

- Cambridge, MA, pages 577–584.
- Nolfi, S. and Floreano, D. (2000). Evolutionary robotics.
- Nolfi, S. and Parisi, D. (1993). Auto-teaching: networks that develop their own teaching input. In *Free University of Brussels*. Citeseer.
- Nolfi, S. and Parisi, D. (1996). Learning to adapt to changing environments in evolving neural networks. *Adaptive behavior*, 5(1):75–98.
- Norouzzadeh, M. S. and Clune, J. (2016). Neuromodulation improves the evolution of forward models. In *Proceedings of the Genetic and Evolutionary Computation Conference*.
- Offerman, T. and Sonnemans, J. (1998). Learning by experience and learning by imitating successful others. *Journal of economic behavior & organization*, 34(4):559–575.
- Oja, E. (1982). A Simplified Neuron Model as a Principal Component Analyzer. *Journal of Mathematical Biology*, 15(3):267–273.
- OpenAI (2017). OpenAI Universe. <https://openai.com/blog/universe/>. Accessed: 21/01/2017.
- Oquab, M., Bottou, L., Laptev, I., and Sivic, J. (2014). Learning and transferring mid-level image representations using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1717–1724.
- Orchard, J. and Wang, L. (2016). The evolution of a generalized neural learning rule. In *Neural Networks (IJCNN), 2016 International Joint Conference on*, pages 4688–4694. IEEE.
- Ormrod, J. E. and Davis, K. M. (2004). *Human learning*. Merrill.
- Paenke, I. (2008). *Dynamics of Evolution and Learning*. PhD thesis, Universität Karlsruhe (TH), Fakultät für Wirtschaftswissenschaften.
- Paenke, I., Kawecki, T., and Sendhoff, B. (2009). The influence of learning on evolution: A mathematical framework. *Artificial Life*, 15(2):227–245.
- Paenke, I., Sendhoff, B., and Kawecki, T. J. (2007). Influence of plasticity and learning on evolution under directional selection. *The American Naturalist*, 170(2):E47–E58.
- Pan, S. J. and Yang, Q. (2010). A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359.
- Pascual-Leone, A., Amedi, A., Fregni, F., and Merabet, L. B. (2005). The plastic human brain cortex. *Annu. Rev. Neurosci.*, 28:377–401.
- Pavlov, I. P. (1927). *Conditioned reflexes*. Oxford : Oxford University Press.
- Pehlevan, C., Hu, T., and Chklovskii, D. B. (2015). A Hebbian/anti-Hebbian neural network for linear subspace learning: A derivation from multidimensional scaling of streaming data. *Neural computation*.
- Pennartz, C. M. A. (1997). Reinforcement Learning by Hebbian Synapses with Adaptive Threshold. *Neuroscience*, 81(2):303–319.
- Pigliucci, M. (2008). Is evolvability evolvable? *Nature Reviews Genetics*, 9(1):75–82.
- Pugh, J., Soltoggio, A., and Stanley, K. O. (2014). Real-time hebbian learning from autoencoder features for control tasks. In *Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*.
- Rauschecker, J. P. and Singer, W. (1981). The effects of early visual experience on the cat’s visual cortex and their possible explanation by Hebb synapses. *Journal of Physiology*, 310:215–239.
- Rawal, A. and Miikkulainen, R. (2016). Evolving deep lstm-based memory networks using an information maximization objective. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016, GECCO ’16*, pages 501–508, New York, NY, USA. ACM.
- Real, E., Moore, S., Selle, A., Saxena, S., Suematsu, Y. L., Le, Q., and Kurakin, A. (2017). Large-Scale Evolution of Image Classifiers. *ArXiv e-prints*.
- Reed, S. and de Freitas, N. (2015). Neural programmer-interpreters. *arXiv preprint arXiv:1511.06279*.
- Rescorla, R. A. (2014). *Pavlovian Second-Order Conditioning (Psychology Revivals): Studies in Associative Learning*. Psychology Press.
- Risi, S., Hughes, C. E., and Stanley, K. O. (2010). Evolving plastic neural networks with novelty search. *Adaptive Behavior*, 18(6):470–491.
- Risi, S. and Stanley, K. O. (2010). Indirectly encoding neural plasticity as a pattern of local rules. In *Proceedings of the 11th International Conference on Simulation of Adaptive Behavior (SAB 2010)*, New York. Springer.
- Risi, S. and Stanley, K. O. (2012). A unified approach to evolving plasticity and neural geometry. In *Neural Networks (IJCNN), The 2012 International Joint Conference on*, pages 1–8. IEEE.
- Risi, S. and Stanley, K. O. (2014). Guided self-organization in indirectly encoded and evolving topographic maps. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO-2014)*. New York, NY: ACM (8 pages).
- Risi, S., Vanderbleek, S. D., Hughes, C. E., and Stanley, K. O. (2009). How novelty search escapes the deceptive trap of learning to learn. In *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages 153–160. ACM.
- Robins, A. (1995). Catastrophic forgetting, rehearsal, and pseudorehearsal. *Connection Science: Journal of Neural Computing, Artificial Intelligence and Cognitive Research*, 7(123-146).
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J. (1988). Learning representations by back-propagating errors. *Cognitive modeling*, 5(3):1.
- Russo, S. J., Dietz, D. M., Dumitriu, D., Morrison, J. H., Malenka, R. C., and Nestler, E. J. (2010). The addicted synapse: mechanisms of synaptic and structural plasticity in nucleus accumbens. *Trends in neurosciences*, 33(6):267–276.
- Rusu, A. A., Rabinowitz, N. C., Desjardins, G., Soyer, H., Kirkpatrick, J., Kavukcuoglu, K., Pascanu, R., and Hadsell, R. (2016). Progressive neural networks. *arXiv preprint arXiv:1606.04671*.
- Sanchez, E., Mange, D., Sipper, M., Tomassini, M., Pérez-Urbe, A., and Stauffer, A. (1996). Phylogeny, ontogeny, and epigenesis: Three sources of biological inspiration for

- softening hardware. In *International Conference on Evolvable Systems*, pages 33–54. Springer.
- Schmidhuber, J. (1991). Curious model-building control systems. In *Neural Networks, 1991. 1991 IEEE International Joint Conference on*, pages 1458–1463. IEEE.
- Schmidhuber, J. (2006). Developmental robotics, optimal artificial curiosity, creativity, music, and the fine arts. *Connection Science*, 18(2):173–187.
- Schmidhuber, J. (2015). Deep learning in neural networks: An overview. *Neural Networks*, 61:85–117.
- Schreiweis, C., Bornschein, U., Burguière, E., Kerimoglu, C., Schreiter, S., Dannemann, M., Goyal, S., Rea, E., French, C. A., Puliadi, R., et al. (2014). Humanized foxp2 accelerates learning by enhancing transitions from declarative to procedural performance. *Proceedings of the National Academy of Sciences*, 111(39):14253–14258.
- Schroll, H. and Hamker, F. H. (2015). Computational models of basal-ganglia pathway functions: focus on functional neuroanatomy. *Frontiers in systems neuroscience*, 7.
- Schultz, W. (1998). Predictive Reward Signal of Dopamine Neurons. *Journal of Neurophysiology*, 80:1–27.
- Schultz, W., Apicella, P., and Ljungberg, T. (1993). Responses of Monkey Dopamine Neurons to Reward and Conditioned Stimuli during Successive Steps of Learning a Delayed Response Task. *The Journal of Neuroscience*, 13:900–913.
- Schultz, W., Dayan, P., and Montague, P. R. (1997). A Neural Substrate for Prediction and Reward. *Science*, 275:1593–1598.
- Silva, F., Urbano, P., and Christensen, A. L. (2012a). Adaptation of robot behaviour through online evolution and neuromodulated learning. In *Ibero-American Conference on Artificial Intelligence*, pages 300–309. Springer.
- Silva, F., Urbano, P., and Christensen, A. L. (2012b). Continuous adaptation of robot behaviour through online evolution and neuromodulated learning. In *Fifth International Workshop on Evolutionary and Reinforcement Learning for Autonomous Robot Systems (ERLARS 2012) in conjunction with the 20th European Conference on Artificial Intelligence (ECAI 2012)*. Montpellier. Citeseer.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., Van Den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., et al. (2016). Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489.
- Simonyan, K. and Zisserman, A. (2014). Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Sims, K. (1991). *Artificial evolution for computer graphics*, volume 25. ACM.
- Sims, K. (1994). Evolving 3d morphology and behavior by competition. *Artificial life*, 1(4):353–372.
- Sipper, M., Sanchez, E., Mange, D., Tomassini, M., Perez-Urbe, A., and Stauffer, A. (1997). A phylogenetic, ontogenetic, and epigenetic view of bio-inspired hardware system. *Evolutionary Computation, IEEE Transactions on*, 1(1):83–97.
- Skinner, B. F. (1938). *The behavior of organisms: An experimental analysis*. New York, London, D. Appleton-Century Company, Incorporated.
- Skinner, B. F. (1953). *Science and Human Behavior*. New York, MacMillan.
- Smith, J. M. (1986). When learning guides evolution. *Nature*, 329(6142):761–762.
- Smith, T. (2002). *The evolvability of artificial neural networks for robot control*. PhD thesis, University of Sussex.
- Smith, T., Husbands, P., Philippides, A., and O’Shea, M. (2002). Neuronal Plasticity and Temporal Adaptivity: Gas-Net Robot Control Networks. *Adaptive Behavior*, 10:161–183.
- Soltoggio, A. (2008a). *Evolutionary and Computational Advantages of Neuromodulated Plasticity*. PhD thesis, School of Computer Science, The University of Birmingham.
- Soltoggio, A. (2008b). Neural Plasticity and Minimal Topologies for Reward-based Learning Problems. In *Proceeding of the 8th International Conference on Hybrid Intelligent Systems (HIS2008), 10-12 September, Barcelona, Spain*.
- Soltoggio, A. (2008c). Neuromodulation Increases Decision Speed in Dynamic Environments. In *Proceedings of the 8th International Conference on Epigenetic Robotics, Southampton, July 2008*.
- Soltoggio, A. (2015). Short-term plasticity as cause-effect hypothesis testing in distal reward learning. *Biological Cybernetics*, 109(1):75–94.
- Soltoggio, A., Bullinaria, J. A., Mattiussi, C., Dürr, P., and Floreano, D. (2008). Evolutionary Advantages of Neuromodulated Plasticity in Dynamic, Reward-based Scenarios. In *Artificial Life XI: Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems*. MIT Press.
- Soltoggio, A., Dürr, P., Mattiussi, C., and Floreano, D. (2007). Evolving Neuromodulatory Topologies for Reinforcement Learning-like Problems. In *Proceedings of the IEEE Congress on Evolutionary Computation, CEC 2007*.
- Soltoggio, A. and Jones, B. (2009). Novelty of behaviour as a basis for the neuro-evolution of operant reward learning. In *Proceedings of the 11th Annual conference on Genetic and evolutionary computation*, pages 169–176. ACM.
- Soltoggio, A. and Stanley, K. O. (2012). From Modulated Hebbian Plasticity to Simple Behavior Learning through Noise and Weight Saturation. *Neural Networks*, 34:28–41.
- Sporns, O. and Alexander, W. H. (2003). Neuromodulation in a learning robot: interactions between neural plasticity and behavior. In *Proceedings of the International Joint Conference on Neural Networks*, volume 4, pages 2789–2794.
- Srivastava, N., Hinton, G. E., Krizhevsky, A., Sutskever, I., and Salakhutdinov, R. (2014). Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15(1):1929–1958.
- Staddon, J. E. R. (1983). *Adaptive Behaviour and Learning*. Cambridge University Press.
- Stanley, K. O. (2007). Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines*, 8(2):131–162.
- Stanley, K. O., D’Ambrosio, D. B., and Gauci, J. (2009). A hypercube-based encoding for evolving large-scale neural



- networks. *Artificial life*, 15(2):185–212.
- Stanley, K. O. and Lehman, J. (2015). Why greatness cannot be planned. *Springer Science Business Media*. doi, 10:978–3.
- Stanley, K. O. and Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2):99–127.
- Stanley, K. O. and Miikkulainen, R. (2003a). Evolving adaptive neural networks with and without adaptive synapses. In Rylander, B., editor, *Genetic and Evolutionary Computation Conference Late Breaking Papers*, pages 275–282, Chicago, USA.
- Stanley, K. O. and Miikkulainen, R. (2003b). A taxonomy for artificial embryogeny. *Artificial Life*, 9(2):93–130.
- Steels, L. (1993). The artificial life roots of artificial intelligence. *Artificial life*, 1(1\_2):75–110.
- Suri, R. E. (2002). TD models of reward predictive responses in dopamine neurons. *Neural Networks*, 15:523–533.
- Suri, R. E., Bargas, J., and Arbib, M. A. (2001). Modeling functions of striatal dopamine modulation in learning and planning. *Neuroscience*, 103(1):65–85.
- Sutton, R. S. and Barto, A. G. (1998). *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA, USA.
- Tessier-Lavigne, M. and Goodman, C. S. (1996). The molecular biology of axon guidance. *Science*, 274(5290):1123.
- Thorndike, E. L. (1911). *Animal Intelligence*. New York: Macmillan.
- Tonelli, P. and Mouret, J.-B. (2013). On the relationships between generative encodings, regularity, and learning abilities when evolving plastic artificial neural networks. *PloS one*, 8(11):e79138.
- Urzelai, J. and Floreano, D. (2000). Evolutionary Robotics: Coping with Environmental Change. In *Genetic and Evolutionary Computation Conference (GECCO’2000)*.
- Vitay, J. and Hamker, F. H. (2010). A computational model of basal ganglia and its role in memory retrieval in rewarded visual memory tasks. *Frontiers in Computational Neuroscience*.
- Wagner, G. P. and Altenberg, L. (1996). Perspective: complex adaptations and the evolution of evolvability. *Evolution*, pages 967–976.
- Walters, E. T. and Byrne, J. H. (1983). Associative Conditioning of Single Sensory Neurons Suggests a Cellular Mechanism for Learning. *Science*, 219:405–408.
- Widrow, B. and Lehr, M. A. (1990). 30 years of adaptive neural networks: perceptron, madaline, and backpropagation. In *Proceedings of the IEEE*.
- Widrow, B., Hoff, M. E., et al. (1960). Adaptive switching circuits. In *IRE WESCON convention record*, volume 4, pages 96–104. New York.
- Widrow, B. and Lehr, M. A. (1990). 30 years of adaptive neural networks: perceptron, madaline, and backpropagation. *Proceedings of the IEEE*, 78(9):1415–1442.
- Yamauchi, B. M. and Beer, R. D. (1994). Sequential behavior and learning in evolved dynamical neural networks. *Adaptive Behavior*, 2(3).
- Yao, X. (1999). Evolving artificial neural networks. *Evolutionary Computation, IEEE Transactions on*, 87(9):1423–1447.
- Yao, X. and Liu, Y. (1997). A new evolutionary system for evolving artificial neural networks. *IEEE transactions on neural networks*, 8(3):694–713.
- Yoder, J. and Yaeger, L. (2014). Evaluating topological models of neuromodulation in polyworld. In *Artificial Life XIV, MIT Press, Cambridge, MA*.
- Young, S. R., Rose, D. C., Karnowski, T. P., Lim, S.-H., and Patton, R. M. (2015). Optimizing deep learning hyper-parameters through an evolutionary algorithm. In *Proceedings of the Workshop on Machine Learning in High-Performance Computing Environments*, page 4. ACM.
- Ziemke, T. and Thieme, M. (2002). Neuromodulation of Reactive Sensorimotor Mappings as Short-Term Memory Mechanism in Delayed Response Tasks. *Adaptive Behavior*, 10:185–199.